

# **VERITAS Volume Replicator 4.1**

## **Administrator's Guide**

**HP-UX**

N13011G

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](http://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

---

<b>Preface</b> .....	<b>xi</b>
Audience .....	xi
How This Guide Is Organized .....	xii
Related VERITAS Documents .....	xiii
Conventions .....	xiv
Getting Help .....	xv
Documentation Feedback .....	xv
<b>Chapter 1. Understanding VERITAS Volume Replicator</b> .....	<b>1</b>
VVR—An Option of VERITAS Volume Manager (VxVM) .....	1
Features of VVR .....	1
What is VVR? .....	2
How VVR Processes Application Writes .....	3
Components of VVR .....	6
Replicated Volume Group .....	6
Storage Replicator Log .....	7
Replication Link—RLINK .....	8
Data Change Map .....	9
Replication Terms Defined in the VVR Context .....	10
Write-Order Fidelity .....	10
Consistent Data Versus Current or Up-to-date Data .....	10
VVR Terms Defined .....	11
Primary and Secondary Host .....	11
Replicated Data Set .....	11



---

Local Host (localhost) .....	12
Heartbeat Protocol .....	12
How the VVR Components Fit Together .....	12
VVR at the Primary .....	12
VVR at the Secondary .....	13
Modes of Replication .....	14
Asynchronous Replication .....	14
Synchronous Replication .....	14
<b>Chapter 2. VVR Concepts .....</b>	<b>17</b>
How Data Flows in VVR Asynchronous Mode .....	17
How Data Flows in VVR Synchronous Mode .....	19
How Data Flows in an RDS Containing Multiple Secondary Hosts .....	21
Understanding How VVR Logs Writes to the SRL .....	23
Understanding Checkpoints .....	25
Understanding Replication Settings for a Secondary .....	27
Mode of Replication—synchronous attribute .....	27
Protection Against SRL Overflow—srlprot attribute .....	28
Latency Protection—latencyprot attribute .....	31
Primary and Secondary Connected .....	32
Primary and Secondary Disconnected .....	32
Understanding the VVR Snapshot Feature .....	33
Traditional Snapshot Feature .....	34
Instant Snapshot Feature .....	35
Instant Full Snapshots .....	35
Instant Space-Optimized Snapshots .....	35
Instant Plex-Breakoff Snapshots .....	36
How VVR Creates Instant Space-Optimized Snapshots .....	37
Comparing the Snapshot Methods Based on Different Features .....	40



---

<b>Chapter 3. Setting Up Replication</b> .....	<b>41</b>
Best Practices for Setting Up Replication .....	42
Creating a Replicated Data Set .....	44
Creating a Primary RVG of an RDS .....	44
Adding a Secondary .....	46
Changing the Replication Settings for a Secondary .....	49
Setting the Mode of Replication .....	50
Setting the Latency Protection .....	51
Setting the SRL Overflow Protection .....	53
Setting the Network Transport Protocol .....	54
Setting the Packet Size .....	54
Controlling the Network Bandwidth Used for Replication .....	55
Synchronizing the Secondary and Starting Replication .....	58
Methods to Synchronize the Secondary .....	58
Using the Network .....	59
Using Block-level Tape Backup .....	60
Moving Disks Physically .....	61
Using the Automatic Synchronization Feature .....	61
Using the Full Synchronization Feature .....	64
Using Block-level Backup and Checkpoint .....	66
Using the Disk Group Split and Join Feature .....	68
Using Difference-Based Synchronization .....	70
Starting Replication When the Data Volumes are Zero Initialized .....	72
Examples for Setting Up a Simple Volume Replicator Configuration .....	73
Creating a Replicated Data Set for the Examples .....	75
Example 1—Setting Up Replication Using Automatic Synchronization .....	76
Example 2—Setting Up Replication Using Full Synchronization .....	76
Example 3—Setting Up Replication Using Block-level Backup and Checkpointing .....	77
Example 4—Setting Up Replication Using Disk Group Split and Join .....	78



---

Example 5—Setting Up Replication Using Differences-Based Synchronization .....	80
Example 6—Setting Up Replication When Data Volumes are Initialized with Zeroes .....	81
<b>Chapter 4. Displaying Configuration Information .....</b>	<b>83</b>
Using the vradmin Display Commands .....	83
Displaying RVG and RDS Information .....	83
Displaying Data Volumes in a Replicated Data Set .....	84
Using the vxprint Display Commands .....	85
Displaying a Specific RLINK .....	85
Interpreting RLINK Flag Settings .....	86
Displaying an Individual RVG .....	88
Interpreting RVG Flag Settings .....	89
Displaying an Individual Data Volume .....	89
Displaying the Status of a Secondary .....	89
Displaying a List of Checkpoints .....	91
Using the vrstat Display Commands .....	92
Displaying the Consolidated Statistics .....	93
Displaying the RLINK Information for All the Hosts in the RDS .....	94
Displaying Information About All the Data Volumes for All the Hosts in the RDS .....	95
Displaying Information About the SRL Volumes for All the Hosts in the RDS .....	96
Displaying Information About the Memory Tunable Parameters for All the Hosts in the RDS .....	97
Determining the Network Bandwidth Being Used by VVR .....	98
Collecting Statistics of the VVR Components .....	99
Understanding How VVR Stores the Statistics .....	101
Displaying Consolidated Replication Status .....	103
Getting and Using Network Performance Data (vxrlink stats) .....	111
Displaying Extended Replication Statistics .....	112





---

Identifying the Most Up-to-Date Secondary .....	114
Example—To Determine the Most Up-to-Date Secondary .....	115
VVR Event Notification .....	117
<b>Chapter 5. Administering VERITAS Volume Replicator .....</b>	<b>121</b>
Administering Data Volumes .....	121
Associating a Volume to a Replicated Data Set .....	121
Verifying the Data on the Primary and Secondary Volumes .....	123
Synchronizing Volumes on the Local Host and Remote Hosts .....	124
Associating the Component Volumes of a Volume Set to an RDS .....	127
Associating a Data Change Map to a Data Volume .....	129
Determining the Region Size .....	130
Resizing a Data Volume in a Replicated Data Set .....	132
Dissociating a Data Volume from its Replicated Data Set .....	134
Administering the SRL .....	135
Protecting from SRL Overflow .....	135
Incrementally Synchronizing the Secondary After SRL Overflow .....	135
Breaking Off Mirrors Before Incremental Synchronization .....	138
Example 1—When snapshot plexes are not available .....	138
Example 2—When snapshot plexes are not available and a disaster occurs .....	139
Resizing the SRL .....	141
Increasing the Size of the SRL on the Primary .....	141
Decreasing the Size of the SRL on the Primary .....	141
Resizing the SRL on the Secondary .....	142
Administering Replication .....	143
Changing the Replication Settings .....	143
Pausing and Resuming Replication to a Secondary .....	143
Stopping Replication to a Secondary .....	144
Changing the IP Addresses Used for Replication .....	145
Changing the Network Ports Used for Replication .....	148



---

Port Numbers Used by VVR .....	148
Displaying and Changing the Ports Used by VVR .....	149
Administering the Replicated Data Set .....	152
Removing a Secondary from a Replicated Data Set .....	152
Removing a Primary RVG .....	153
Administering Checkpoints .....	154
Creating Checkpoints .....	154
Ending Checkpoints .....	155
Viewing Checkpoints .....	155
Deleting Checkpoints .....	155
Creating RVG Snapshots .....	157
Using the Instant Snapshot Feature .....	157
Instant Full Snapshot .....	158
Creating Snapshot Volumes for Data Volumes in an RVG .....	159
Preparing the Volumes .....	159
Freezing or Pausing Replication .....	160
Creating Instant Full Snapshots .....	161
Unfreezing or Resuming Replication .....	163
Instant Space-Optimized Snapshots .....	163
Preparing the RVG Volumes for Snapshot Operation .....	164
Creating the Cache Object .....	165
Freezing or Pausing Replication .....	165
Creating Instant Space-Optimized Snapshots .....	166
Unfreezing or Resuming Replication .....	167
Instant Plex-breakoff Snapshots .....	168
Preparing the RVG Volumes for Snapshot Operation .....	168
Creating Snapshot Plexes for Data Volumes in an RVG .....	168
Freezing or Pausing Replication .....	168
Creating Instant Plex Breakoff Snapshots .....	169
Unfreezing or Resuming Replication .....	170



---

Administering Snapshots .....	171
Refreshing Snapshots .....	171
Reattaching the Snapshot Plexes to the Data Volumes (snapback) .....	171
Restoring Data from the Snapshots .....	173
Displaying the Snapshot Information .....	175
Destroying the Snapshots .....	176
Using the Traditional Snapshot Feature .....	177
Creating Snapshot Plexes for Data Volumes in an RVG .....	177
Freezing or Pausing Replication .....	178
Taking a Snapshot .....	178
Unfreezing or Resuming Replication .....	178
Performing the Required Operations on the Snapshot .....	179
Reattaching the Snapshot Plexes to the Data Volumes (Snapback) .....	179
Using Snapback with resyncfromreplica Option .....	180
Using VERITAS Volume Manager FastResync .....	182
Enabling FastResync .....	182
Verifying the DR Readiness of a VVR Setup .....	183
Performing a Failover .....	183
Performing a Firedrill .....	184
Verifying the Data on the Secondary .....	185
Performing Online Data Verification .....	185
Performing Offline Data Verification .....	187
Backing Up the Secondary .....	191
Checkpoint Pause/Resume Secondary RLINK .....	191
Restoring the Secondary from Online Backup .....	192
Restoring from Secondary Checkpoints .....	192
Restoring a Secondary RLINK .....	193
Changing the VVR Tunables .....	194
Changing the Tunable Values Using vxtune .....	195
Changing the VVR Tunables Using SAM .....	197



---

<b>Chapter 6. Using VVR for Off-Host Processing</b> .....	<b>199</b>
What is Off-Host Processing? .....	199
In-Band Control Messaging Overview .....	199
How to Use the Data on the Secondary .....	200
Using Snapshots .....	200
In-Band Control Messaging Explained .....	201
Performing Off-Host Processing Tasks .....	204
Tasks to Perform for Off-Host Processing .....	205
Using the IBC Messaging Command <code>vradmin ibc</code> .....	206
Prerequisites for Using <code>vradmin ibc</code> Command .....	206
Understanding the Scripts Used for the <code>vradmin ibc</code> Command .....	208
Sample Scripts .....	210
Examples of Off-Host Processing .....	211
Example 1—Decision Support Using the Snapshot Feature and the <code>vradmin ibc</code> Command .....	212
Example 2—Backing Up Using the Snapshot Feature and the <code>vradmin ibc</code> Command .....	213
Example 3—Performing Block-Level Backup of the Secondary Data Using the <code>vradmin ibc</code> Command .....	214
 <b>Chapter 7. Transferring the Primary Role</b> .....	 <b>215</b>
Migrating the Primary .....	216
Important Notes for Migrating the Primary Role .....	218
Example 1—Migrating from a Healthy Primary .....	219
Example 2—Migrating the Primary Role in a Setup with Multiple Secondaries .....	220
Taking Over from an Original Primary .....	223
Important Notes About Taking Over from an Original Primary .....	225
About Fast Failback .....	226
Example 1—Taking Over from an Original Primary .....	228
Example 2—Taking Over from an Original Primary in a Setup with Multiple Secondaries .....	229



Failing Back to the Original Primary .....	231
Fast Failback Versus Difference-Based Synchronization .....	231
Failing Back Using Fast Failback Synchronization .....	232
Example 1—Failing Back to the Original Primary Using Fast Failback .....	234
Example 2—Failing Back to the Original Primary Using Fast Failback in a Multiple Secondaries Setup .....	236
Failing Back Using Difference-Based Synchronization .....	237
Converting an Original Primary to a Secondary .....	237
Example 3—Failing Back to the Original Primary Using Difference-Based Synchronization .....	238
Example 4—Failing Back to the Original Primary Using Difference-Based Synchronization in a Multiple Secondaries Setup .....	239
<b>Chapter 8. Administering VVR Using VVR VEA .....</b>	<b>241</b>
Features of VVR VEA .....	242
Replicated Data Set Defined .....	243
Getting Started with VVR VEA .....	244
Launching VEA .....	244
From a UNIX Client .....	244
From a Windows Client .....	244
Connecting to the Hosts .....	245
Closing VVR .....	245
Working with VVR .....	245
VVR VEA Main Window .....	246
Tree and Detailed Views .....	246
Replication Menu on the Menu Bar .....	248
Pop-up Menus .....	248
Setting Up Replication Using VVR VEA .....	249
Overview of Setting Up Replication .....	249
Example Scenario .....	251
Setting Up Replication Using the Create RDS Wizard .....	252



---

Setting Up Replication Using Individual Tasks .....	258
Creating a Replicated Data Set .....	258
Synchronizing the Secondary and Starting Replication .....	260
Start Replication Dialog Box .....	261
Starting Replication Using Automatic Synchronization .....	261
Setting Up Replication Using Checkpoint .....	262
Starting Replication When Data Volumes are Initialized with Zeroes .....	264
Viewing Replication Information .....	265
Viewing the RDSs Present on the Selected Host .....	265
Using the RDS List View .....	265
Viewing RDS Information .....	267
Using the RDS View .....	267
Viewing Information About RVGs .....	269
Using the RVG Views .....	270
Viewing Information About the Primary RVG .....	270
Viewing Information About the Secondary RVG .....	271
Administering Data Volumes .....	274
Associating a Data Volume to an RDS .....	274
Resizing a Data Volume in a Replicated Data Set .....	275
Dissociating a Data Volume from its RDS .....	276
Increasing the Size of the SRL on the Primary .....	277
Administering Replication .....	278
Changing the Replication Settings .....	278
Pausing Replication to a Secondary .....	279
Resuming Replication to a Secondary .....	280
Stopping Replication to a Secondary .....	280
Changing the Replication Network for a Secondary .....	282
Administering the Replicated Data Set .....	284
Removing a Secondary from an RDS .....	284
Removing a Primary RVG .....	285



---

Incrementally Synchronizing the Secondary After SRL Overflow .....	285
Migrating the Primary Role .....	287
Taking Over from an Original Primary .....	289
Failback Options .....	289
Failing Back to the Original Primary .....	292
Failing Back Using Fast-Failback Synchronization .....	293
Example—Failing Back to the Original Primary Using Fast Failback .....	294
Failing Back Using Automatic Synchronization .....	295
Converting a Primary to a Secondary .....	296
Synchronize the Secondary RVG with the Primary RVG .....	297
Example: Failing Back Using Automatic Synchronization .....	297
Event Notification Using VVR VEA .....	299
Accessing the Rule Manager Utility .....	299
Preliminary Setup .....	300
Main Steps for Creating Rules .....	302
Accessing the VERITAS Volume Replicator SNMP MIB Files .....	302
<b>Chapter 9. Troubleshooting VVR .....</b>	<b>303</b>
Recovery from RLINK Connect Problems .....	303
Recovery from Configuration Errors .....	306
Errors During an RLINK Attach .....	306
Errors During Modification of an RVG .....	309
Missing Data Volume Error .....	309
Data Volume Mismatch Error .....	310
Data Volume Name Mismatch Error .....	311
Recovery .....	312
Primary-Host Crash .....	312
Recovering from Primary Data Volume Error .....	312
Primary SRL Volume Error Cleanup and Restart .....	314
RVG PASSTHRU Mode .....	315



---

Primary SRL Volume Error at Reboot .....	315
Primary SRL Volume Overflow Recovery .....	316
Primary SRL Header Error Cleanup and Recovery .....	316
Recovering from SRL Header Error .....	317
Secondary Data Volume Error Cleanup and Recovery .....	317
Recovery Using a Secondary Checkpoint .....	318
Cleanup Using a Primary Checkpoint .....	318
Secondary SRL Volume Error Cleanup and Recovery .....	319
Secondary SRL Header Error Cleanup and Recovery .....	320
Secondary SRL Header Error at Reboot .....	322
<b>Appendix A. VVR Command Reference .....</b>	<b>323</b>
Command Reference .....	323
<b>Appendix B. Messages .....</b>	<b>333</b>
Kernel Messages .....	333
Error Messages .....	333
Messages Related to RLINKs .....	333
Messages Related to the SRL and DCM .....	338
Messages Related to Communication Errors .....	341
Messages Related to Configuration Errors .....	345
Message Related to I/O Failure .....	346
Messages Related to Kernel Logging .....	347
Informational Messages .....	349
Informational RLINK Messages .....	349
Informational SRL and DCM Messages .....	351
Utility Error Messages .....	352
The vradmin Error Messages .....	354
Messages Related to Configuration Errors .....	364
Messages Related to the vrstat Command .....	370





---

<b>Appendix C. Using the In-Band Control Messaging Utility vxibc and the IBC Programming API</b>	<b>371</b>
In-Band Control Messaging Overview	372
Using the IBC Messaging Command-Line Utility	373
Registering an Application Name	373
Displaying the Registered Application Name for an RVG	373
Receiving an IBC Message	374
Sending an IBC Message	374
Unfreezing the Secondary RVG	375
Unregistering an Application Name	375
Receiving and Processing an IBC Message Using a Single Command	376
Sending and Processing an IBC Message Using a Single Command	377
Examples—Off-Host Processing	378
Example 1—Decision Support Using the Traditional Snapshot Feature and the vxibc Utility	380
Example 2—Backing Up Using the Snapshot Feature and the vxibc Utility	381
Example 3—Trial Failover Using the Snapshot Feature	382
Example 4—Decision Support Using the Instant Full Snapshot Feature and the vxibc Utility	383
In-Band Control Messaging API	384
IOCTL Commands	385
RV_IBC_REGISTER	385
RV_IBC_SEND	387
RV_IBC_RECEIVE	388
RV_IBC_UNFREEZE	390
RV_IBC_UNREGISTER	390
Using the IBC API	391
<b>Appendix D. VERITAS Volume Replicator Object States</b>	<b>393</b>
VERITAS Volume Replicator Kernel State	393
RVG KSTATES	393



RLINK KSTATES .....	393
VERITAS Volume Replicator Utility States .....	395
RVG Utility States .....	395
RLINK Utility States .....	395
Inactive RLINKs .....	396
STALE RLINK State .....	396
FAIL RLINK State .....	397
Inconsistent RLINKs .....	398
Pausing, Resuming, and Restoring RLINK States .....	399
<b>Appendix E. VVR Task Reference .....</b>	<b>401</b>
Mapping the Name of a Secondary Data Volume to a Differently Named Primary Data Volume .....	401
Mapping Disk Groups .....	403
Decreasing the Size of the SRL on the Primary .....	404
<b>Appendix F. Accessibility and VVR VEA .....</b>	<b>407</b>
Keyboard Navigation and Shortcuts in VVR VEA .....	408
VEA Exceptions to Microsoft Keyboard Conventions .....	408
VEA Additions to Microsoft Keyboard Conventions .....	409
VEA Help Additions to Microsoft Keyboard Conventions .....	411
General Keyboard Navigation Within the GUI .....	411
Keyboard Navigation Within Dialog Boxes .....	412
Support for Accessibility Settings and Assistive Technologies .....	415
<b>Glossary .....</b>	<b>417</b>
<b>Index .....</b>	<b>427</b>



# Preface

---

The *VERITAS™ Volume Replicator Administrator's Guide* provides information on configuring, using, and administering VERITAS Volume Replicator (VVR). This guide introduces VVR concepts and explains how to administer VVR from the command-line and the menu interface.

## Audience

This guide is for system administrators who configure and maintain VVR and assumes:

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



## How This Guide Is Organized

[Chapter 1, “Understanding VERITAS Volume Replicator” on page 1](#) gives an overview of VERITAS Volume Replicator, introduces VVR components and the important concepts of VVR.

[Chapter 2, “VVR Concepts” on page 17](#) explains the important concepts of VVR and thus prepares you to use VVR effectively.

[Chapter 3, “Setting Up Replication” on page 41](#) describes the commands used to configure and administer VVR, and gives the steps to set up a simple VVR configuration using the command-line interface.

[Chapter 4, “Displaying Configuration Information” on page 83](#) provides the commands to display a consolidated view of the Replicated Data Sets (RDSs) in a VVR environment.

[Chapter 5, “Administering VERITAS Volume Replicator” on page 121](#) describes how to perform RVG tasks and manage Secondaries. It also explains how to set up replication modes, latency protection, and SRL protection and describes how to synchronize the Secondary.

[Chapter 6, “Using VVR for Off-Host Processing” on page 199](#) describes how to use VVR to perform off-host processing tasks.

[Chapter 7, “Transferring the Primary Role” on page 215](#) describes how to transfer the Primary role using the VERITAS Volume Replicator (VVR) command-line interface.

[Chapter 8, “Administering VVR Using VVR VEA” on page 241](#) describes how to configure, monitor, and administer VVR in a distributed environment.

[Chapter 9, “Troubleshooting VVR” on page 303](#) explains how to recover from VVR configuration and replication errors.

[Appendix A, “VVR Command Reference” on page 323](#) lists all of the VVR commands and a description of each command.

[Appendix B, “Messages” on page 333](#) contains a list of error messages generated by VVR, accompanied by a definition of each message and a suggested user action if applicable.

[Appendix C, “Using the In-Band Control Messaging Utility vxibc and the IBC Programming API” on page 371](#) explains how to use the individual commands of the In-Band Control (IBC) Messaging Utility vxibc and the IBC Messaging Application Programming Interface (API).

[Appendix D, “VERITAS Volume Replicator Object States” on page 393](#) describes the Kernel State (KSTATE) and utility state (STATE) of the RVG and RLINK objects.

[Appendix E, “VVR Task Reference” on page 401](#) lists some additional administrative tasks.



[Appendix F, “Accessibility and VVR VEA” on page 407](#) provides information about the accessibility features and compliance of VERITAS Enterprise Administrator for VERITAS Volume Replicator (VVR VEA).

[“Glossary”](#) on page 417 contains a list of terms and definitions relevant to VVR.

## 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.
<b>monospace (bold)</b>	Indicates user input.	<b># ls pubs</b> <b>C:\&gt; dir pubs</b>
<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.
<b>bold</b>	Depicts GUI objects, such as fields, list boxes, menu selections, etc. Also depicts GUI commands.	Enter your password in the <b>Password</b> field. Press <b>Return</b> .
<a href="#">blue text</a>	Indicates hypertext links.	See " <a href="#">Getting Help</a> " on page xv.
#	Unix superuser prompt (all shells).	<b># cp /pubs/4.0/user_book /release_mgnt/4.0/archive</b>
C:\>	Windows user prompt.	<b>C:\&gt; copy \pubs\4.0\user_book c:\release_mgnt\4.0\archive</b>



## 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](mailto: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>.







# Understanding VERITAS Volume Replicator

---

1

VERITAS Volume Replicator (VVR) is data-replication software designed to contribute to an effective disaster recovery plan. VVR enables you to maintain an exact or consistent copy of application data at one or more remote locations.

VVR replicates writes from the source host to the remote hosts. If a disaster occurs at the source location, you can use the copy of the application data at the remote location to restart the application at the remote location. The host at the source location on which the application is running is known as the Primary host, and the host at the target location is known as the Secondary host. You can have up to 32 Secondary hosts in a VVR environment.

## VVR—An Option of VERITAS Volume Manager (VxVM)

VVR is an optionally licensed feature of VERITAS Volume Manager (VxVM) that works as a fully integrated component of VxVM. VVR benefits from the robustness, ease of use, and high performance of VxVM, and at the same time, adds replication capability to VxVM. VVR can replicate existing VxVM configurations, and can be transparently configured while the application is active.

## Features of VVR

The features of VERITAS Volume Replicator (VVR) include:

- ◆ Replicates data for up to 32 remote locations over any IP network in a LAN or WAN environment.
- ◆ Can be configured to work with any storage hardware supported by VERITAS Volume Manager.
- ◆ Provides volume-level replication of application or file system data, which includes support for all commercial database management systems, such as Oracle, DB2, Sybase, and Informix.
- ◆ Easily accommodates growth of application data and system configurations.



- ◆ Supports replication of VxVM ISP (Intelligent Service Provisioning) volumes in an RVG, provided all the volumes are of the same type.
- ◆ Performs replication of volume groups in asynchronous or synchronous modes, ensuring data integrity and consistency in both modes.
- ◆ Maintains write-order fidelity, which applies writes on the Secondary host in the same order that they were issued on the Primary host.
- ◆ Provides the command-line interface and the graphical interface for online management of the VVR environment.
- ◆ Provides multiple methods to synchronize the data at the Secondary location with the data at the Primary location.
- ◆ Enables you to easily recover your application at the remote site.
- ◆ Supports off-host processing such as Decision Support Systems (DSS) and backup, by enabling you to break off a consistent mirror or snapshot of the data volumes on the Secondary to use for these operations.

## What is VVR?

VVR is an extension of the logical volume management capability of VxVM and is included with VxVM. You can start using VVR in a VxVM configuration by installing a valid VVR license.

VVR replicates initially synchronized volumes from a source location to one or more remote locations across any distance. It provides an exact or consistent copy of application data at the remote locations. VVR provides several methods to initialize the application data between the primary location and the remote location, such as using the network, using tape backup, and moving disks physically.

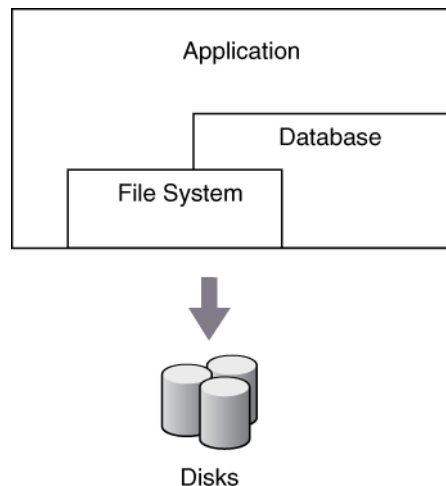
## How VVR Processes Application Writes

This section helps you understand how application writes are directed when VxVM is not being used, when VxVM is added, and when VVR is added.

### When VxVM/VVR is not being used

When VxVM is not being used, the application writes to a file system placed on a disk partition. In the case of applications or databases on raw devices, the database writes directly to the disk partition, instead of to a file system. In either case, the application, that is, the database or a file system, sends data to the operating system to be written to disk and the operating system communicates directly with the disks.

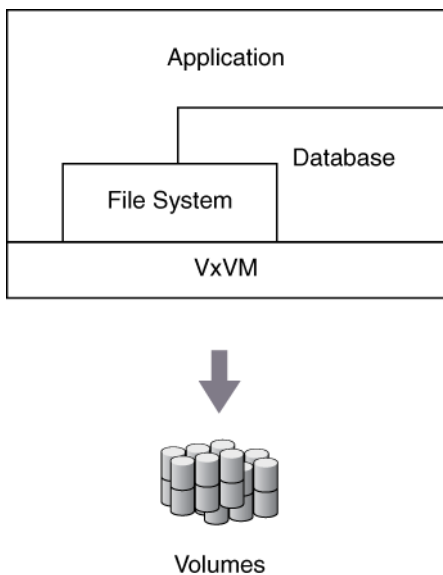
How Application Writes are Processed When VxVM and VVR are not Used



### When VxVM is present

When VxVM is being used, the applications write to logical devices called volumes, which have been created on Volume Manager disks (VM disks) rather than physical disks. A volume is a virtual disk device that appears to applications, such as databases, and file systems as a physical disk. However, a volume does not have the limitations of a physical disk.


How Application Writes are Processed When VxVM is Used



## When VVR is enabled

When VVR is added, it resides between the application and the underlying VxVM volumes. All writes to these replicated volumes are intercepted and replicated to the Secondary host in the order in which they were received at the Primary. Writes are also applied to the local volumes. However, reads are directly processed using the local volumes.

How Application Writes are Processed When VxVM and VVR are Used



VVR sends writes to the Secondary in the order in which they were received on the Primary. The Secondary receives writes from the Primary and writes to local volumes.

The data volumes on the Secondary cannot be used while replication is in progress: that is, you cannot run an application directly on the data volumes on the Secondary. If an application on the Secondary attempts to access the data volumes, it will see inconsistent data because the application on the Primary is constantly changing the data on the volumes. To use the data on the Secondary while the Primary is active, use the snapshot feature to create a version of the data which is not being changed.

If a disaster occurs on the Primary, the applications that were running on the Primary can be brought up on the Secondary.

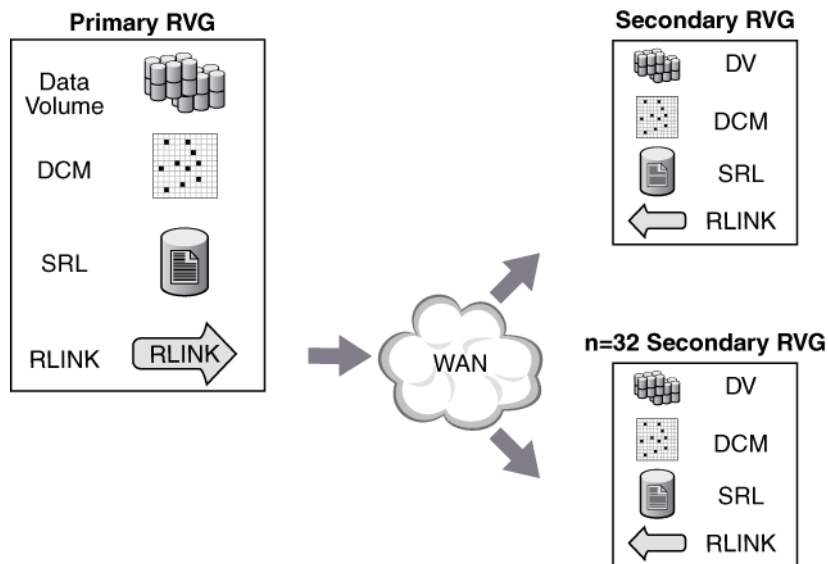


## Components of VVR

This section explains the components of VVR, which contain configuration information. The components include:

- ◆ [Replicated Volume Group \(RVG\)](#)
- ◆ [Storage Replicator Log \(SRL\)](#)
- ◆ [Replication Link—RLINK](#)
- ◆ [Data Change Map \(DCM\)](#)

Sample Configuration to Illustrate the Different VVR Components



## Replicated Volume Group

In the case of a database, several processes perform writes to disks. Database processes write in a specific order. This order must be maintained at all times including when recovering from a disk failure. For example, the database posts any database change to the log before writing to the table space. To convey to VVR that these two volumes are related, these two volumes must be grouped.

A Replicated Volume Group (RVG) is a group of volumes within a given VxVM disk group configured for replication. An RVG is always a subset of a VxVM disk group. One or more related volumes in a disk group can be configured as an RVG. By related volumes, we mean a set of volumes to which application writes must be replicated in order on the Secondary.

All related volumes must be part of the same disk group. Unrelated volumes must not be grouped together in an RVG. Multiple RVGs can be configured inside one disk group, although this is not a recommended configuration.

Volumes that are associated with an RVG and contain application data are called *data volumes*. The data volumes in the RVG are under the control of an application, such as a Database Management System, that requires write-order fidelity among the writes to the volumes.

Write-ordering is strictly maintained within an RVG during replication to ensure that each remote volume is always consistent, both internally and with all other volumes of the group. Each RVG can have a maximum of 1023 data volumes. VVR replicates data from a Primary RVG to a Secondary RVG.

The concept of Primary and Secondary is per RVG, not per system. A system can simultaneously be a Primary host for some RVGs and Secondary host for others. An RVG also contains the Storage Replicator Log (SRL) and Replication Link (RLINK), which are explained in the following sections in [“Storage Replicator Log”](#) on page 7 and [“Replication Link—RLINK”](#) on page 8.

---

**Note** A Primary RVG can have multiple Secondary RVGs. When this document refers to the Secondary host, it implicitly refers to all the Secondary RVGs.

---

## Storage Replicator Log

The Storage Replicator Log (SRL) is a circular buffer of writes for an RVG. Each RVG contains one SRL. Writes to the data volumes in the RVG are first queued in the SRL on the Primary host before they are sent to the Secondary. VVR uses the SRL to track the order of writes to data volumes in the RVG. The SRL enables VVR to maintain write-order fidelity at the Secondary RVG.

From a VxVM perspective, the SRL is just another volume. Because all writes are written to the SRL first, it is important for the SRL to have optimum write performance. This means all performance techniques used to increase write performance of a volume apply to the SRL. For most implementations, the SRL is striped across multiple drives for write performance, and mirrored to an equal set of drives for protection.



Each write to the disks generates two writes: one to the SRL, and one to a data volume. For this reason, the data volumes and SRL volumes must be configured on different physical disks to improve write performance. Note that VVR does not allow application writes to the SRL.

The SRL on the Secondary host performs a different function from the one on the Primary. Under normal operations, the Secondary SRL is not used. While VVR is recovering after a temporary failure in communication between the Primary and Secondary, or from a Primary or Secondary host failure, certain writes are stored in the Secondary SRL and applied together to maintain data consistency. This ensures that the Secondary transitions from one consistent data state to another.

## Replication Link—RLINK

An RLINK is associated with an RVG and establishes the link between the Primary and a Secondary RVG. Each RLINK associated with a Primary RVG represents one Secondary. Each RLINK associated with a Secondary RVG represents a Primary. The attributes of an RLINK specify the replication parameters for the corresponding Secondary. For example, the network to be used for replication between the Primary and the Secondary can be specified as the attribute of the RLINK.

A Primary RVG can have up to 32 associated RLINKs. Although a Secondary RVG can also have 32 associated RLINKs, it can have only one active RLINK; this active RLINK represents the Primary that is currently replicating to this Secondary RVG.

VVR reads data from the SRL and sends it to each Secondary. Each Secondary receives data at its own rate. A write is removed from the SRL when all the Secondary RVGs have successfully received the writes. If a Secondary does not keep up with the write rate, the SRL could overflow for the corresponding RLINK.



## Data Change Map

VVR uses the Data Change Map (DCM) to track writes when the SRL overflows and thus enables you to avoid complete resynchronization of the data on the Secondary. DCM is a component of VVR that contains a bitmap and can be optionally associated with a data volume on the Primary RVG. It is almost identical to the DRL in a VxVM volume, but the bits represent ranges that are different between the Primary and the Secondary, rather than between mirrors.

While the DCM is active, each bit that has been set in the DCM represents a region whose contents are different between the Primary and the Secondary. The DCM becomes active only when the SRL is no longer large enough to hold accumulated updates. At an appropriate time, the administrator initiates a resynchronization and causes VVR to incrementally synchronize the Secondary with the Primary by looking up the bitmap. Note that the Secondary is inconsistent between the time the DCM resynchronization starts and completes because write-order fidelity is not preserved. After the resynchronization is complete, the Secondary RVG is consistent and replication resumes as usual.

The Automatic Synchronization, SRL Overflow Protection with DCM, and Fast Failback features use the DCM. Each data volume in the RVG must have a valid DCM associated with it before the DCM can be used.

Under normal circumstances, writes are sent to the Secondary in the order in which they are received at the Primary. Consequently, the Secondary data is consistent with the Primary and can be used to take over the Primary if a disaster occurs. Note that under normal circumstances, the DCM is not in use.

When the resynchronization of the DCM starts, the Secondary becomes inconsistent because the DCM resynchronization writes are not necessarily in the same order as the application writes. The Secondary becomes consistent only after the resynchronization of the DCM is complete. As a result, the Secondary cannot be used for disaster recovery while the DCM resynchronization continues.



## Replication Terms Defined in the VVR Context

This section defines the following replication terms in the VVR context:

- ◆ [Write-Order Fidelity](#)
- ◆ [Consistent Data Versus Current or Up-to-date Data](#)

### Write-Order Fidelity

To use the Secondary in a disaster recovery scenario, write-order fidelity must be maintained. The term *write-order fidelity* means that VVR tracks writes on the Primary in the order in which they are received and applies them on the Secondary in the same order. It is important to maintain write-order fidelity to ensure that the data on the Secondary is consistent with the data on the Primary. While the data at the Secondary can be behind in time, it must be a consistent image of the Primary RVG at a known point in the past.

Without write order fidelity, there is no guarantee that a Secondary has consistent, recoverable data. VVR maintains write-order fidelity regardless of the mode of replication and across all the data volumes in an RVG. For example, in a database environment, the log and data space are typically on different volumes. On the Primary, VVR applies writes to the log and data spaces in a fixed order and maintains this fixed order when applying the writes on the Secondary. If write-order fidelity is not maintained, a database application may not recover successfully when failed over to the Secondary.

### Consistent Data Versus Current or Up-to-date Data

Data is *consistent* if the system or application using it can be successfully restarted to a known, usable state. The data on the Secondary is consistent if it correctly reflects the data on the Primary at some point in the past. At all times, VVR maintains the data at the Secondary in a consistent state with the data at the Primary. For example, if the data being replicated is used by a database, the data is consistent if the database can be started and recovered to a usable state with zero data corruption. If the data contains a file system, the data is consistent if the file system check utility can be run and it can recover with no file system corruption.

Data is considered consistent only if it contains all updates up to a point in time and none of the updates that come after that point. For example, if it is a file system, the most recently created files may be missing when it is restarted. Or, if it is a database, one or more of the most recently committed transactions might be missing.

Data that is *current* or *up-to-date* contains the latest changes made at the Primary. For example, if you are replicating a database, the most recent transaction is available at the Secondary. Whether or not the data on the Secondary must always be current is a business decision and can be controlled by choosing between synchronous and asynchronous modes of replication. Synchronous mode guarantees that the data on the Secondary is current, at the cost of application performance; asynchronous mode does not guarantee the data is current, but has less impact on application performance and provides the ability to use more cost-effective telecommunications.

## VVR Terms Defined

This section defines the following VVR terms:

- ◆ [Primary and Secondary Host](#)
- ◆ [Replicated Data Set](#)
- ◆ [Local Host \(localhost\)](#)
- ◆ [Heartbeat Protocol](#)

### Primary and Secondary Host

The concept of Primary host and Secondary host are used only in the context of a particular Replicated Data Set (RDS), not per system. A system in the configuration can simultaneously be the Primary for one RDS and Secondary for others. This allows for very flexible replication configurations.

### Replicated Data Set

VVR replicates data from a Primary host, where the application is running, to one or more Secondary hosts. A Replicated Volume Group (RVG) on the Primary host and its counterparts on the Secondary hosts make up an RDS. An RDS is not a Volume Manager object but just a concept used in VVR. An RDS enables grouping of the RVG on the Primary and its counterparts on the Secondaries.

Most VVR commands operate on an RDS, that is, the Primary RVG and all the Secondaries in the RDS. You can issue VVR commands from any host in an RDS unless otherwise noted. VVR performs the appropriate tasks on the required hosts in the RDS.



## Local Host (localhost)

The host from which a command is issued is called the *local host*. The name of the Replicated Volume Group (RVG) on the local host represents the RDS. For example, to add a data volume to an RDS, issue the command from any host in the RDS, using the name of the RVG on that host to specify the RDS; VVR adds a data volume to the corresponding RVGs on all the hosts in the RDS.

## Heartbeat Protocol

To ensure that the Secondary host can always detect communication loss regardless of update activity, the Primary host periodically sends a heartbeat message to the Secondary. If the Secondary misses a fixed number of heartbeat messages, it detects a communication loss and begins its recovery protocol. Likewise, if the Primary is unable to send a heartbeat message or if its heartbeat messages go unacknowledged, the RVG on the Primary host also detects a communication loss and enters its recovery procedure.

On successful completion of the recovery protocol, update activity resumes automatically unless some interim administrative command or error prevents it. Heartbeat messages use the UDP protocol. The choice of protocol is available for data and control messages.

## How the VVR Components Fit Together

This section describes how the VVR components fit together to enable replication.

- ◆ [VVR at the Primary](#)
- ◆ [VVR at the Secondary](#)

## VVR at the Primary

VVR is configured such that the volumes to be replicated for a specific application are placed in an RVG. Writes to the data volumes are persistently queued in the SRL. The SRL on the Primary tracks all writes in the order in which they were received and transmits the writes to the Secondary using a replication link (RLINK). The SRL is a VxVM volume configured as part of an RVG. The SRL gives VVR the ability to associate writes to specific volumes within the Replicated Data Set in a specific order, and thus to maintain write-order fidelity at the Secondary RVG.

On the Primary, each write to an RVG generates two writes: one to a data volume, and one to the SRL. The write to the data volume is written in the background and does not affect application performance. While VVR generates two writes, only the write to the SRL affects the application.

The write to the SRL is a fast write to a sequentially accessed log while the data volume write is a normal write performed asynchronously. The write to the data volume is not in the critical path for the application.

If the Primary crashes at any point before the write completes to the data volume, data is fully recoverable from the SRL. This is very similar to a database writing to a redo log and later writing to the data files. This two phase write gives VVR the ability to maintain write-order fidelity at the Secondary.

## VVR at the Secondary

VVR sends data to the Secondary RVG as a message encompassing an application write. This means VVR sends messages based on application write sizes. When the Secondary receives the message in VVR kernel memory, the Secondary immediately sends an initial acknowledgement of receipt. This is known as the *network acknowledgement*. The network acknowledgement allows the Primary to immediately continue processing, as required. The data is not yet written to disk on the Secondary RVG, but it is still safe because it is stored in the Primary SRL. After the Secondary writes to the local disk, it sends the second acknowledgement, the *data acknowledgement*. The write is discarded from the SRL when the Primary receives the data acknowledgement.

The reason for the two-phase acknowledgement is performance. In synchronous mode, the Primary waits for the network acknowledgement from the Secondary before it completes the write at the application. If VVR were to wait for the write to complete on the Primary and the Secondary, it would increase latency considerably. By using the two-phase acknowledgement, VVR maintains application performance. Because data is persistently queued in the Primary SRL, safety of the data for the Secondary is maintained.

VVR receives a packet into memory on the Secondary RVG when using UDP, holds the packet until the all previous packets have been received, then writes to the disks in the correct sequence to maintain consistency at the Secondary. Holding the packets in memory enables VVR to reassemble out-of-order network traffic before writing, and discover and handle missing packets. To maintain consistency at the Secondary RVG, VVR never writes an I/O out of order with the Primary RVG. VVR serializes and checksums incoming data from the Primary RVG to support accurate replay to the Secondary volumes.

During normal data transfer the Secondary SRL is not in the data path. The Secondary SRL is only used in very specific conditions, as follows:

- ◆ During recovery after a Primary or Secondary crash
- ◆ To store state of actual underlying volume plexes
- ◆ When an IBC freezes a Secondary



## Modes of Replication

VVR replicates in *synchronous* and *asynchronous* modes. You can set the mode of replication between the Primary and each Secondary depending on your requirement. You can replicate to the Secondary hosts of an RDS in different replication modes.

In each mode, VVR replicates and completes the application writes differently. Each mode deals differently with network conditions. The following sections provide a brief description of synchronous and asynchronous modes of replication and discuss some of the issues involved in choosing between the two. For more information, see the *VERITAS Volume Replicator Planning and Tuning Guide*.

### Asynchronous Replication

Asynchronous mode is useful when it is acceptable for the Secondary not to be up-to-date. When replicating in asynchronous mode, an update to the Primary volume is complete when it has been recorded in the Primary SRL. In asynchronous mode, all completed updates to the Primary volumes are guaranteed to be made on the Secondary data volumes with some delay. This is true despite failures in communication or system crashes on any of the participating hosts.

The application is informed that the write request is complete and the write is queued persistently to be sent to the Secondary. This queue may grow when there is a surge in the write rate. The queue is being continuously drained. When the surge subsides, the queue drains faster than it grows enabling the Secondary to catch up with the Primary. Because asynchronous mode queues writes persistently and holds them at the Primary for later transmission, it is able to deal with temporary outages of the network or the Secondary host without affecting the performance of the application. If a disaster strikes in asynchronous mode, it is likely that the most recent writes have not reached the Secondary. Asynchronous mode has the disadvantage that the data at a Secondary might not have the latest updates when a failover occurs. For more information about asynchronous mode, see “Asynchronous Mode Considerations” in the *VERITAS Volume Replicator Planning and Tuning Guide*.

### Synchronous Replication

Synchronous mode ensures that a write has been posted to the Secondary and the Primary before the write completes at the application level. In synchronous mode, the data on the Secondary is completely up-to-date and if a disaster occurs at the Primary, data can be recovered from any surviving Secondary without any loss. If the Secondary must reflect all writes that have successfully completed on the Primary, synchronous mode is the correct choice.

Synchronous replication keeps the Secondary up-to-date with the Primary by waiting for each write to reach the Secondary before the application sees the successful completion of the write on the Primary.

Synchronous replication provides data currency but can impact application performance in high latency or limited bandwidth environments. When using synchronous replication, the response time experienced by the application is affected because the write has to wait for the Secondary to acknowledge it before the write can complete on the Primary.

Synchronous replication is most effective in application environments with low update rates. However, it can be deployed in write-intensive environments where high bandwidth, low latency network connections are available.

The performance of synchronous replication could degrade significantly if the peak application write rate exceeds the available network bandwidth. The amount of degradation can be reduced by increasing network bandwidth and reducing network latency between the Primary and Secondary. Refer to the *VERITAS Volume Replicator Planning and Tuning Guide* for a discussion of network latency and network bandwidth, and their effects on VVR performance.

In synchronous mode, the `synchronous` attribute enables you to specify what action is taken when the Secondary is unreachable. The `synchronous` attribute can be set to `override` or `fail`. When the `synchronous` attribute is set to `override`, synchronous mode converts to asynchronous during a temporary outage. In this case, after the outage passes and the Secondary catches up, replication reverts to synchronous.

When the `synchronous` attribute is set to `fail`, the application receives a failure for writes issued while the Secondary is unreachable. The application is likely to fail or become unavailable, and hence this setting must be chosen only if such a failure is preferable to the Secondary being out of date.

We recommend setting the `synchronous` attribute to `override`, as this behavior is suitable for most applications. Setting the `synchronous` attribute to `fail` is suitable only for a special class of applications that cannot have even a single write difference between the Primary and Secondary data volumes. In other words, this mode of operation must be used only if you want an application write to fail if the write cannot be replicated immediately. It is imperative that the network connection between hosts using this option must be highly reliable to avert unnecessary application downtime as network outage could cause an application outage. For more information about synchronous mode, see “Synchronous Mode Considerations” in the *VERITAS Volume Replicator Planning and Tuning Guide*.







# VVR Concepts

---

2

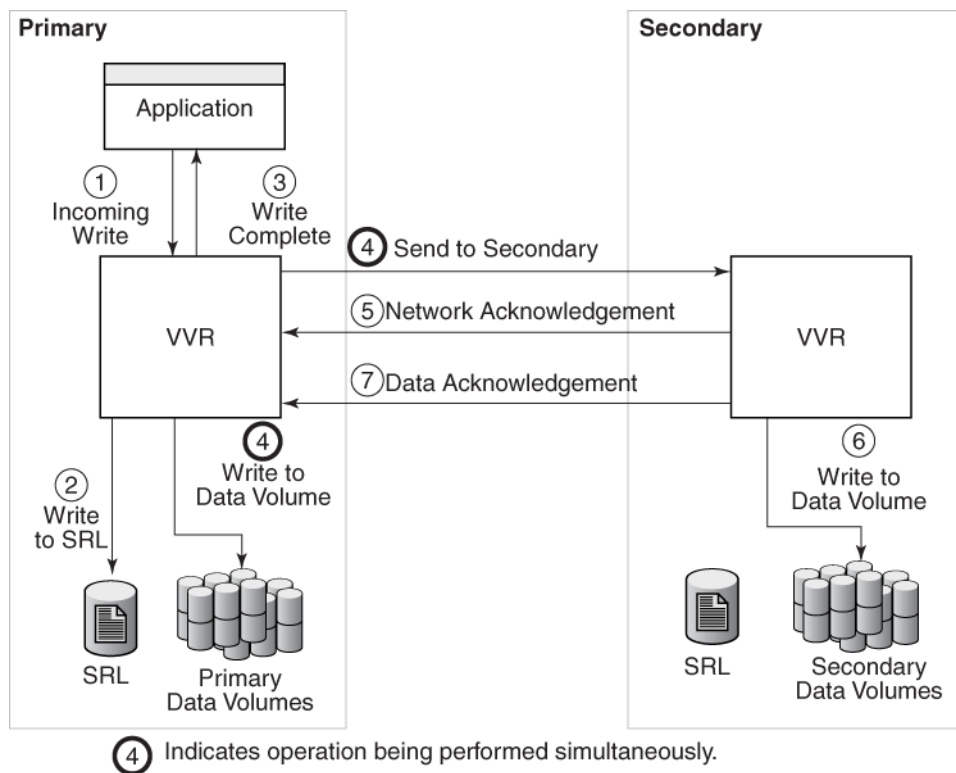
Understanding the VERITAS Volume Replicator (VVR) concepts is the key to getting the maximum out of VVR. This chapter explains the important concepts of VVR and thus prepares you to use VVR effectively. Read this chapter before setting up replication.

## How Data Flows in VVR Asynchronous Mode

This section explains how VVR processes an incoming write when replicating in asynchronous mode.



Example—How Data Flows in the Asynchronous Mode of Replication

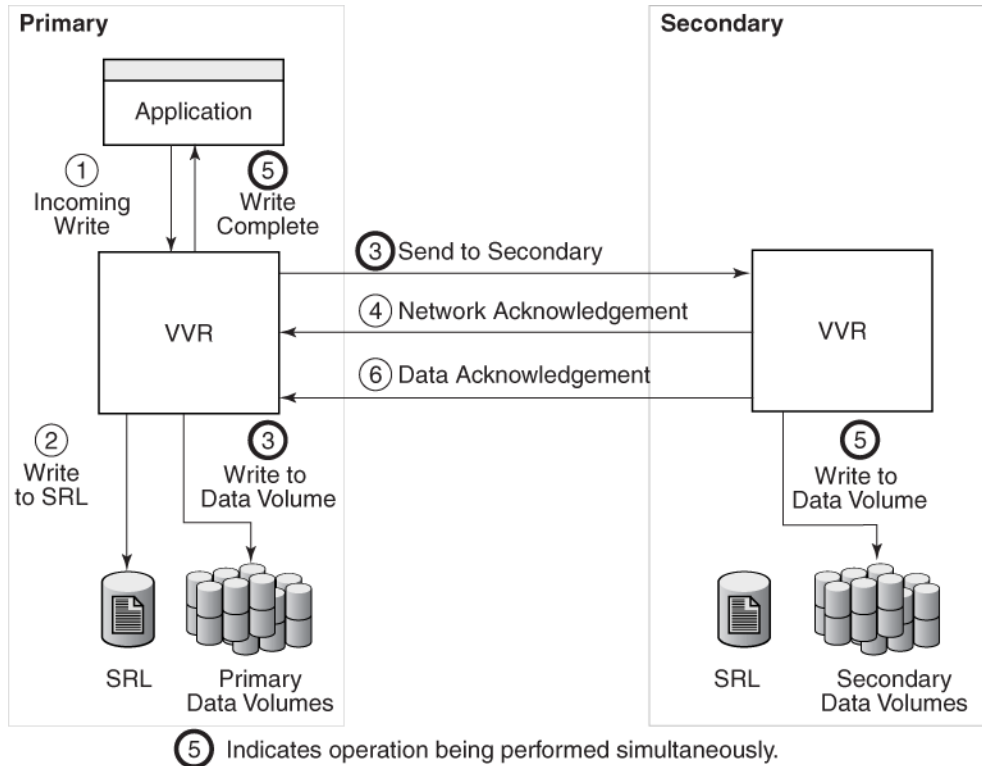


1. VVR receives a write on the Primary.
2. Writes it to the Primary SRL.
3. On the Primary, acknowledges to the application that the write is complete.
4. Sends the writes to the asynchronous Secondary hosts, in the order in which they were received on the Primary, and at the same time, writes to the Primary data volumes.
5. When the Primary receives the network acknowledgement, it knows that the write has been received in the Secondary VVR memory buffer.
6. VVR sends the writes to the data volumes on the Secondary and then sends a data acknowledgement to the Primary.
7. When the Primary receives the data acknowledgement, VVR marks the write as complete in the SRL.

## How Data Flows in VVR Synchronous Mode

This section explains how VVR processes an incoming write when replicating in synchronous mode.

Example—How Data Flows in the Synchronous Mode of Replication



In synchronous mode of replication, VVR processes an incoming write as follows:

1. VVR receives a write on the Primary.
2. Writes it to the Primary SRL.
3. Sends the write to the Secondary hosts and waits for the synchronous network acknowledgements from the Secondary hosts. At the same time, VVR writes to the data volumes on the Primary.
4. On the Secondary, VVR receives the write, processes it, and sends a network acknowledgement to the Primary.



5. Sends writes to the data volumes on the Secondary; when the Primary receives a network acknowledgement from all the Secondary hosts, VVR acknowledges to the application that the write is complete.

Note that the Secondary RVG sends the network acknowledgement as soon as the write is received in the VVR kernel memory. This removes the time required to write to the Secondary data volumes from the application latency. On the Primary, VVR does not wait for data to be written to the Secondary data volumes. This improves application performance. However, VVR tracks all such acknowledged writes that have not been written to the data volumes. VVR can replay these tracked writes if the Secondary crashes before writing to the data volumes on the Secondary or if the Primary crashes before it receives the data acknowledgement.

6. When the write is written to the data volumes on the Secondary, VVR on the Secondary sends a data acknowledgement to the Primary. VVR marks the write as complete in the SRL when the Primary receives the data acknowledgement from all the Secondary hosts.

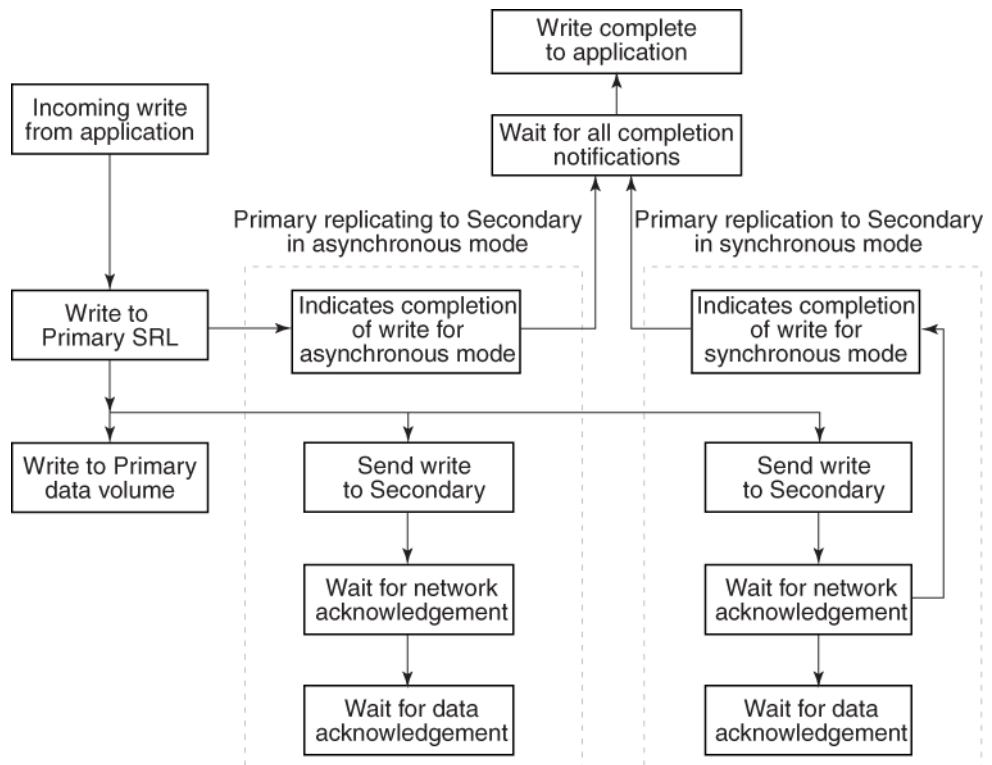
When an RDS containing multiple Secondary RVGs is replicating in synchronous mode, the application latency is determined by the slowest Secondary. Overall performance in synchronous mode is determined by the time to write to the SRL, plus the round-trip time required to send data to the Secondary RVG and receive the acknowledgement.



## How Data Flows in an RDS Containing Multiple Secondary Hosts

This section explains how VVR processes an incoming write for a Replicated Data Set (RDS) containing multiple Secondary hosts, some replicating in asynchronous mode and some in synchronous mode.

How Data Flows in an RDS with Multiple Secondaries



In asynchronous and synchronous mode of replication, VVR processes an incoming write as follows:

1. Receives a write from the application.
2. Writes it to the SRL.
3. Sends the write to the Secondary hosts replicating in synchronous mode and in asynchronous mode. At the same time, VVR writes to the data volume on the Primary.



4. On the Secondary, VVR receives the write, processes it, and sends a network acknowledgement to the Primary.
5. When the Primary receives a network acknowledgement from the Secondary hosts replicating in synchronous mode, VVR acknowledges to the application that the write is complete.

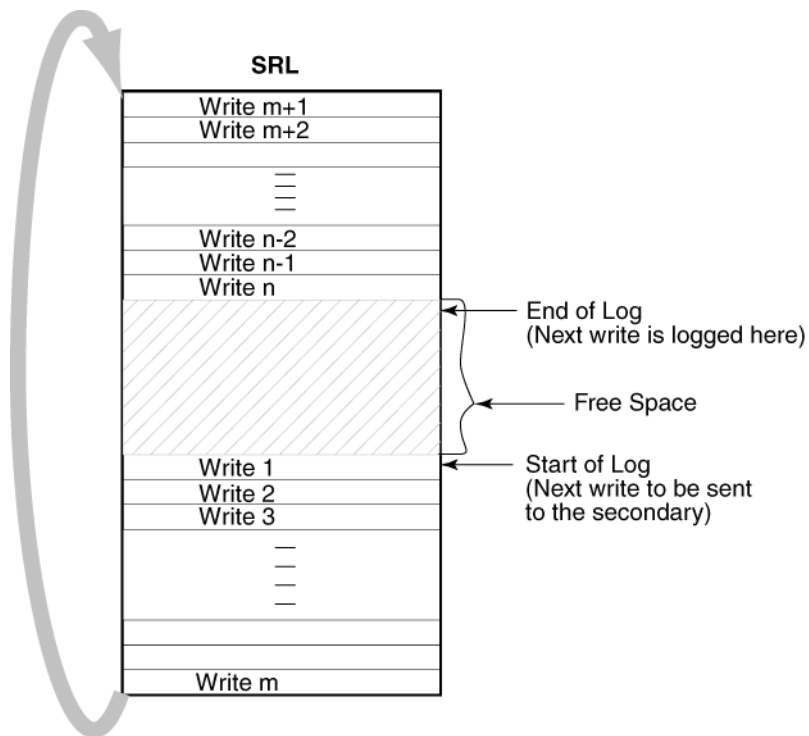
Note that the Secondary RVG sends the network acknowledgement as soon as the write is received in the VVR kernel memory. This removes the time required to write to the Secondary data volumes from the application latency. On the Primary, VVR waits only for the network acknowledgement from all the synchronous Secondary hosts and not for the data to be written to the Secondary data volumes. This improves application performance. However, VVR tracks all such acknowledged writes that have not been written to the data volumes. VVR can replay these tracked writes if the Secondary crashes before writing to the data volumes on the Secondary or if the Primary crashes before receiving the data acknowledgement.

6. When the write is written to the data volumes on the Secondary, VVR sends a data acknowledgement from the Secondary to the Primary in both synchronous and asynchronous mode.
7. When the Primary receives the data acknowledgement from all the Secondary hosts, VVR marks the write as complete in the SRL.

## Understanding How VVR Logs Writes to the SRL

VVR receives writes from the application and queues them in the SRL for transmission to the Secondary hosts. All the RLINKs of an RVG share the SRL. The SRL header contains a specific set of pointers for each RLINK that indicates the writes that have not been sent to the corresponding Secondary.

This section explains the working of the SRL as a circular buffer.



The illustration shows how writes are logged in the SRL. The earliest write that came in is Write 1, which also represents the Start of Log for the Secondary. VVR logs Write 2, Write 3, Write m one after the other until it reaches the end of the SRL. Because the SRL is a circular log the next write, Write m+1 wraps around and logging continues. When the Primary receives the data acknowledgement from this Secondary host for Write 1, VVR marks the Write 1 as complete in the SRL. VVR then processes Write 2, Write 3, and so on.

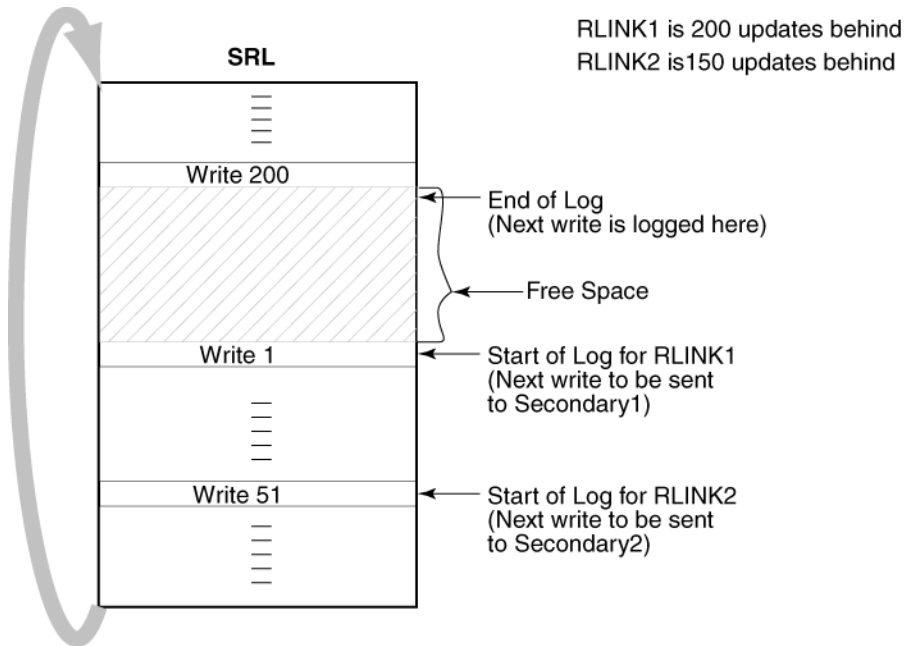
VVR maintains the following types of pointers in the SRL header:

- ◆ The *Start of Log* pointer—Each RLINK has a Start of Log pointer that designates the next write, Write 1, to be sent to the Secondary.



- ◆ The *End of Log* pointer—Designates the location to be written to by the next incoming write after Write n.

In the following illustration, RLINK1 is 200 writes or updates behind, whereas RLINK2 is 150 writes behind. If the End of Log pointer reaches the Start of Log pointer of the RLINK, the SRL overflows for this RLINK.



Synchronous RLINKs are usually up-to-date. Typically, the Start of Log and End of Log pointers of synchronous RLINKs are separated by the number of simultaneous I/O operations the application performs. For asynchronous RLINKs, the difference between the Start of Log pointer and End of Log pointers reflect how many outstanding writes have yet to be processed, that is, how behind is the RLINK. Different RLINKs usually have Start of Log pointers indicating different places in the SRL; this reflects the difference in the rate at which data is sent to the Secondary.



## Understanding Checkpoints

VVR checkpoints are user-defined markers in the SRL. Each checkpoint has a start (checkstart) and an end (checkend). Checkpoints are used to perform the following tasks:

- ◆ Synchronizing the Secondary while the Primary application is active
- ◆ Restoring the Secondary data volumes

The Secondary data volumes must be synchronized with the Primary data volumes before replication can start: that is, after adding a Secondary to the RDS, after a Secondary data volume error, or after SRL overflow. VVR enables you to synchronize the Secondary data volumes while the application is active on the Primary. If you use the automatic synchronization feature of VVR to synchronize the Secondary data volumes over the network, VVR ensures that the Secondary data volumes are consistent and up-to-date when the synchronization process completes. However, you can also synchronize the Secondary data volumes by performing a backup of the Primary data volumes and applying it on Secondary or by copying the data over the network using the VVR `vradmin` command or any other utility. If the Primary application is active during the synchronization process, the Secondary data volumes are inconsistent and not up-to-date when the synchronization is complete.

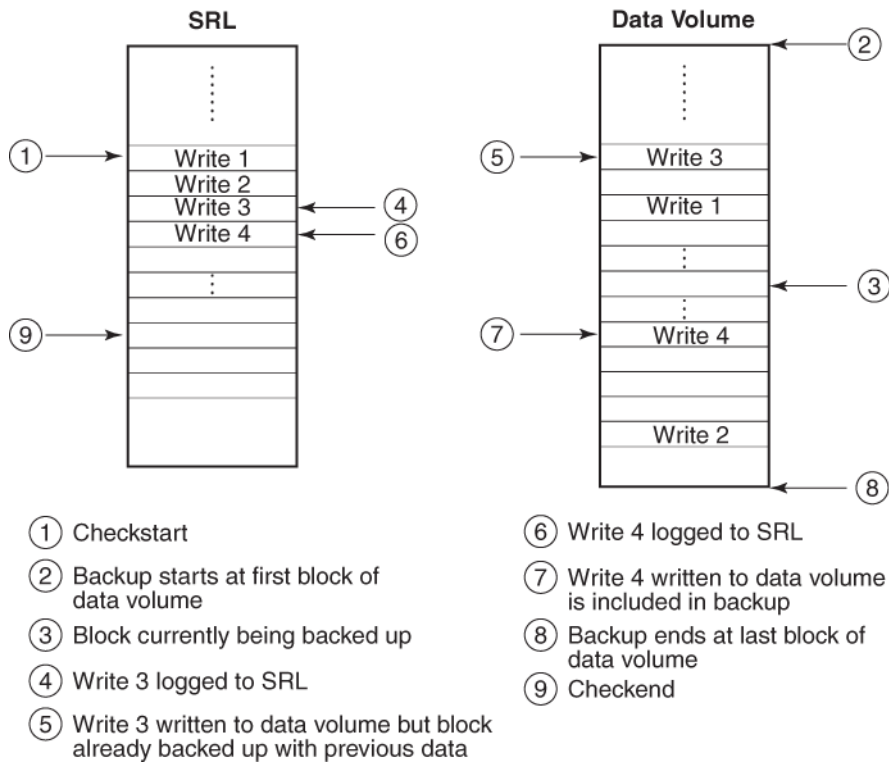
Typically, a backup or synchronization utility performs sequential reads starting with the first block of the volume until it reaches the last block of the volume and transfers those blocks either to tape or over the network. If the Primary application is active during this process, some Primary data volume blocks might have changed while the data volumes are read sequentially. It is likely that the application changes several blocks, some of which are read by the synchronization process before they were changed and some after they were changed. This results in the Secondary data volumes being inconsistent and not completely up-to-date at the end of the synchronization process.

To make the Secondary consistent and up-to-date, VVR must transfer in order all the blocks that changed during the synchronization process. In a VVR environment, all writes to the Primary data volumes are logged to the SRL; therefore, VVR can transfer the writes that occurred during the synchronization to the Secondary. To do this, VVR must know the start and end of the synchronization process. VVR checkpoints are used to indicate this start position (checkstart) and end position (checkend) in the SRL.

Because the checkpoint information is stored in the SRL, checkpoints become invalid when the SRL wraps around. The same checkpoint and tape backups can be used to synchronize the data volumes on multiple Secondary hosts if the checkpoint remains valid.

VVR enables you to create a maximum of forty-six checkpoints. If the number of checkpoints exceeds this number VVR displays an error message asking you to delete the earlier checkpoints. You can selectively delete the required checkpoints.





As shown in the illustration in this section, a backup utility may copy previous contents of the blocks corresponding to Write 3 (event 5) but copy updated contents of block corresponding to Write 4 (event 7). However, as shown in the illustration, VVR logs all the writes to the SRL (events 4 and 6). Note that a checkstart was performed (event 1) before the backup was started (event 2) and a checkend was performed (event 9) after the backup was completed (event 8). On starting replication with this checkpoint after the synchronization is complete, VVR can transfer all the writes between checkstart and checkend and make the Secondary data volumes up-to-date and consistent.

## Understanding Replication Settings for a Secondary

The VVR replication settings determine the replication behavior between the Primary RVG and a specific Secondary RVG. VVR behaves differently based on the settings for mode of replication, SRL overflow protection, and latency protection, depending on whether the Secondary is connected or disconnected. To use the replication settings effectively in your environment, it is important to understand how each replication setting affects replication when the Primary and Secondary are connected and disconnected. A Secondary is said to be disconnected from the Primary if the RLINK becomes inactive because of a network outage or administrative action.

VVR enables you to set up the replication mode, latency protection, and SRL protection using the three replication attributes `synchronous`, `latencyprot`, and `srlprot`. These attributes are of the form `attribute=value`. Each attribute setting could affect replication and must be set up with care.

### Mode of Replication—synchronous attribute

VVR replicates in two modes: synchronous and asynchronous. In synchronous mode, a write must be posted to the Secondary and the Primary before the write completes at the application level. When replicating in asynchronous mode, an update to the Primary volume is complete when it has been recorded in the Primary SRL. The decision to use synchronous or asynchronous mode must be made with an understanding of the effects of this choice on the replication process and the application performance.

You can set up VVR to replicate to a Secondary in synchronous or asynchronous mode by setting the `synchronous` attribute of the RLINK to `override`, `off`, or `fail`. The following table summarizes how the state of the RLINK affects the mode of replication:

Value of <code>synchronous</code> Attribute	When RLINK is Connected	When RLINK is Disconnected
<code>override</code>	Synchronous	Asynchronous
<code>off</code>	Asynchronous	Asynchronous
<code>fail</code>	Synchronous	I/O error to application



`synchronous=off`

By default, VVR sets the `synchronous` attribute to `off`. Setting the attribute of an RLINK to `synchronous=off` sets the replication between the Primary and the Secondary to asynchronous mode.

`synchronous=override`

Setting the `synchronous` attribute to `override` puts the RLINK in synchronous mode and specifies override behavior if the RLINK is disconnected. During normal operation, VVR replicates in synchronous mode. If the RLINK is disconnected, VVR switches temporarily to asynchronous mode and continues to receive writes from the application and logs them in the SRL. After the connection is restored and the RLINK is up-to-date, the RLINK automatically switches back to synchronous mode.

`synchronous=fail`

---

**Caution** You must read the section “Synchronous Mode Considerations” in the *VERITAS Volume Replicator Planning and Tuning Guide* if you use the `synchronous=fail` mode.

---

Setting the `synchronous` attribute to `fail` puts the RLINK in synchronous mode and specifies the behavior if the RLINK is disconnected. During normal operation, VVR replicates in synchronous mode. If the RLINK is disconnected, VVR fails incoming writes to the Primary.

## Protection Against SRL Overflow—`srlprot` attribute

If the network is down or the Secondary is unavailable, the number of writes in the SRL waiting to be sent to the Secondary could increase until the SRL fills up. When the SRL cannot accommodate a new write without overwriting an existing one, the condition is called SRL overflow. At this point, the new writes are held up or the RLINK overflows depending on the mode of SRL overflow protection.

The circumstances that could cause the SRL to overflow include:

- ◆ A temporary burst of writes or a temporary congestion in the network causing the current update rate to exceed the currently available bandwidth between the Primary and the Secondary.
- ◆ A temporary failure of the Secondary node or the network connection between the Secondary and the Primary.
- ◆ Replication is paused by an administrator.
- ◆ The network bandwidth is unable, on a sustained basis, to keep up with the update rate at the Primary. This is not a temporary condition and can be corrected only by increasing the network bandwidth or reducing the application update rate, if possible.

If the SRL overflows, the Secondary becomes out-of-date and must be completely synchronized to bring it up-to-date with the Primary. The SRL Protection feature of VVR enables you to either prevent the SRL from overflowing or tracks the writes using the Data Change Map (DCM) if the SRL overflows. You must weigh the trade-off between allowing the overflow or affecting the application. You can prevent SRL overflow using the `srlprot` setting.

VVR provides the following modes of SRL overflow protection: `autodcm`, `dcm`, or `override`. VVR activates these modes only when the SRL overflows. You can set up SRL protection by setting the `srlprot` attribute of the corresponding RLINKs to `autodcm`, `dcm`, or `override`. By default, the `srlprot` attribute is set to `autodcm`. The following table summarizes how the state of the RLINK affects SRL protection when the SRL is about to overflow:

Value of the <code>srlprot</code> Attribute	When RLINK is Connected	When RLINK is Disconnected
<code>autodcm</code>	Convert to DCM logging	Convert to DCM logging
<code>dcm</code>	Protect*	Convert to DCM logging
<code>override</code>	Protect*	Overflow

\* Protect by stalling application writes until SRL drains 5% to 95% full or drains 20 megabytes, whichever is smaller.

If the SRL overflow protection is set to `autodcm`, `override`, or `dcm`, SRL overflow protection is enabled. The replication setting for the Secondary and the state of the connection between the Primary and the Secondary determines how VVR works when the SRL is about to overflow.

`srlprot=autodcm`

VVR activates the DCM irrespective of whether the Primary and Secondary are connected or disconnected. Each data volume in the RVG must have a DCM; note that VVR does not stall writes when `srlprot` is set to `autodcm`.

`srlprot=dcm`

If the Primary and Secondary are connected, new writes are stalled in the operating system of the Primary host until a predetermined amount of space, that is 5% or 20 MB, whichever is less, becomes available in the SRL.

If the Primary and Secondary are disconnected, DCM protection is activated and writes are written to the DCM. Each data volume in the RVG must have a DCM.



### `srlprot=override`

If the Primary and Secondary are connected, new writes are stalled in the operating system of the Primary host until a predetermined amount of space, that is 5% or 20 MB, whichever is less, becomes available in the SRL.

If the Primary and Secondary are disconnected, VVR disables SRL protection and lets the SRL overflow.

For more information, see [“Setting the SRL Overflow Protection”](#) on page 53.



## Latency Protection—`latencyprot` attribute

Excessive lag between the Primary and the Secondary could be a liability in asynchronous replication. The Latency Protection feature of VVR protects the Secondary host from falling too far behind in updating its copy of data when replicating in asynchronous mode. This feature limits the number of outstanding writes lost in a disaster enabling automatic control of excessive lag between Primary and Secondary hosts when you replicate in asynchronous mode.

When replicating in asynchronous mode, it is normal for the SRL to have writes waiting to be sent to the Secondary. If your network has been sized based on the average update rate of the application on the Primary node, the number of writes waiting in the Primary SRL is likely to be within an acceptable range. The number of writes in the SRL would grow under the following circumstances:

- ◆ A temporary burst of writes or a temporary congestion in the network, which causes the current update rate to exceed the currently available bandwidth between the Primary and the Secondary.
- ◆ A temporary failure of the Secondary node or the network connection between the Secondary and the Primary.
- ◆ Replication is paused by an administrator.
- ◆ The network bandwidth is unable, on a sustained basis, to keep up with the write rate at the Primary. This is not a temporary condition and can be corrected only by increasing the network bandwidth or reducing the application write rate if possible.

If the Primary SRL has a large number of writes waiting to be transferred to the Secondary, the Secondary data is considerably behind the Primary. If a disaster strikes the Primary and the Secondary takes over, the Secondary would not contain all the data in the Primary SRL. In this case, the data on the Secondary would be consistent but significantly out of date when the Secondary takes over. To prevent the Secondary from being too far behind the Primary in this scenario, you can limit the number of writes waiting in the Primary SRL for transmission to the Secondary by setting up latency protection.

Latency protection has two components, its mode, and the `latency_high_mark` and `latency_low_mark` which specify when the protection is active or inactive. The `latency_high_mark` specifies the maximum number of waiting updates in the SRL before the protection becomes active and writes stall or fail, depending on the mode of latency protection.

The `latency_low_mark` must be a number lower than the `latency_high_mark`; the `latency_low_mark` is the number of writes in the SRL when the protection becomes inactive and writes succeed. You can set up latency protection by setting the `latencyprot` attribute to either `override` or `fail`. Setting the attribute to `latencyprot=off`, which is the default, disables latency protection.



Setting the attribute to `latencyprot=fail` or `override` enables latency protection. The following sections explain how VVR controls replication depending on the setting of the `latencyprot` attribute of the RLINK when the Primary and Secondary either connected or disconnected.

The following table summarizes how the state of the RLINK affects the latency protection:

Value of <code>latencyprot</code> Attribute	When RLINK is Connected	When RLINK is Disconnected
<code>override</code>	Protect	Drop protection
<code>off</code>	No protection	No protection
<code>fail</code>	Protect	I/O error to application

## Primary and Secondary Connected

`latencyprot=fail` or `override`

Under normal operation, if the number of waiting writes increase and reach the `latency_high_mark`, following writes are stalled in the operating system of the Primary until the SRL drains sufficiently to bring the number of waiting writes below the `latency_low_mark`.

## Primary and Secondary Disconnected

Primary and Secondary are said to be disconnected when they are in the PAUSED state or are disconnected because of a network outage, or an outage of the Secondary node.

`latencyprot=override`

VVR allows the number of writes in the SRL to exceed the `latency_high_mark`. In this case, VVR causes latency protection to be overridden and allows incoming writes from the application whose data is being replicated. VVR does not stall incoming writes because the SRL is currently not draining and incoming writes may be stalled indefinitely. Stalling of incoming writes is undesirable for the writing application. Most system administrators set `latencyprot=override`.



If replication is paused and not resumed, or if there is an extended network outage, the outstanding writes can exceed the latency high mark. When the Secondary reconnects either because replication is resumed or the network becomes available, VVR starts stalling writes until the writes in the SRL reach the latency low mark. The time required for the Primary to send the accumulated writes to the Secondary can take a long time depending on the amount of data to be sent and the bandwidth of the network. The application perceives this as VVR being unresponsive and some applications may time out resulting in application error.

`latencyprot=fail`

If the number of writes in the SRL reaches the `latency_high_mark` while the Primary and the Secondary are disconnected, VVR causes new writes at the Primary to fail. This prevents the Secondary from falling further behind than specified by the `latency_high_mark`.

## Understanding the VVR Snapshot Feature

VVR enables you to create an image of the online data volumes, at a given point in time and such an image is referred to as a snapshot. The data in the original volume may change; however, the snapshot can still be used as a stable and independent copy for various purposes, including the following:

- ◆ to restore data both on the Primary and Secondary if the original data gets corrupted because of logical errors, administrative errors, or media errors such as disk failures.
- ◆ to verify the Disaster Recovery (DR) readiness of the DR site or perform firedrill activities.
- ◆ to create a copy of the data for application development or testing.
- ◆ to support off-host processing for applications such as Decision Support Systems (DSS) or for report generation.
- ◆ to perform online data verification of the volumes in an RVG when replication is in progress.
- ◆ to retain a consistent copy of the Secondary data volumes during Data Change Map (DCM) resynchronization.

---

**Note** You can use the snapshot feature on the Primary and the Secondary host.

---

The snapshot feature in VVR is the same as the snapshot feature in VxVM, because VVR is fully integrated with VxVM. In VVR, an RVG is a group of VxVM volumes; therefore, taking a snapshot of an RVG is the same as taking a snapshot of the data volumes in the RVG.



VVR provides two methods of creating snapshots which are:

- ◆ [Traditional Snapshot Feature](#)
- ◆ [Instant Snapshot Feature](#)

## Traditional Snapshot Feature

The traditional snapshot feature of VVR enables you to create snapshots of all the data volumes in an RVG at a single point in time, by breaking off the volume plexes. You can create snapshots when the volume plexes are completely synchronized with the data volume.

This method requires you to create and attach the appropriate snapshot plexes that are the same size as the original volumes, before taking the snapshot. For more information on creating the plexes, refer to the *VERITAS Volume Manager Administrator's Guide*.

After creating and attaching the snapshot plexes, they must be synchronized with the data volume. The time required for synchronizing the plexes in the beginning is directly proportional to the size of the volume. Thus, depending on the size of the volume it may take a long time before the plexes are ready to be used.

When the plexes are synchronized, you can then create snapshots after freezing replication using the IBC commands or after pausing replication.

For more information on traditional snapshots, refer to the section "[Using the Traditional Snapshot Feature](#)" on page 177.

## Instant Snapshot Feature

The instant snapshot feature enables you to take instant full snapshots, instant space-optimized snapshots, or instant plex-breakoff snapshots. The advantages of this method compared to the traditional method are:

- ◆ the plexes or snapshot volumes do not require synchronization before taking the snapshots.
- ◆ the snapshots are instantly available.

The instant snapshot feature provided by VVR can be used either from the Primary or the Secondary. VVR also provides you with the option to take space-optimized snapshots.

The types of instant snapshots that you can create are:

- ◆ [Instant Full Snapshots](#)
- ◆ [Instant Space-Optimized Snapshots](#)
- ◆ [Instant Plex-Breakoff Snapshots](#)

---

**Note** If an RVG consists of only VxVM ISP volumes, you can perform only the instant full snapshot operation and the instant space-optimized snapshot operations. When creating the instant space-optimized snapshot, you must create the VxVM ISP cache object before proceeding with the operation as VVR does not support the *cache\_size* parameter for an RVG with VxVM ISP volumes.

---

### Instant Full Snapshots

The instant full snapshot feature of VVR enables you to create a full snapshot of all the data volumes in an RVG without any delay. In this case, the snapshot plexes do not require synchronization before creating the snapshot. Therefore, the required data is available instantly after the snapshot is created. However, this method requires the snapshot volumes to be created with the appropriate naming convention, before you proceed with the snapshots. For more information on creating instant full snapshots, refer to the section “[Instant Full Snapshot](#)” on page 158.

### Instant Space-Optimized Snapshots

VVR also enables you to take instant space-optimized snapshots. Unlike instant full snapshots, the instant space-optimized snapshots require less storage space than the original volume because space-optimized snapshots store only the data that has changed between the original volume and the snapshot. Typically, the data that has changed between the original volume and the snapshot volume over the lifetime of the snapshot is minimal compared to the total data on the volume. Thus, the space-optimization achieved can be significant. The snapshot data is managed by VVR using a cache object which



functions as a space-optimized persistent store. You must create the cache object before you take the instant space-optimized snapshots or specify the size of the cache object. Multiple snapshots can be created on the same cache object. The cache object can be created with an *autogrow* option set to *on* to allow it to grow automatically if the cache volume size is not enough for the specified writes. For more information on the cache object, refer to the section “[Preparing the RVG Volumes for Snapshot Operation](#)” on page 164. For more information on creating instant space-optimized snapshots, refer to the section “[Creating Instant Space-Optimized Snapshots](#)” on page 166.

## Instant Plex-Breakoff Snapshots

Similar to the traditional plex-breakoff snapshot feature, this method also requires the plexes to be attached to the source volume before taking the snapshots. Although the synchronization of the plexes may still take a long time, the major difference between the traditional snapshots and the instant plex-breakoff snapshots is that you can instantly start performing the operations such as refresh, restore, and snapback on the instant plex-breakoff snapshots.

The instant plex-breakoff snapshots operation requires the plexes to be named using the `plexprefix` attribute if you want to use specific plexes. Otherwise, VVR uses the plexes that are in the `SNAPDONE` state. For more information on creating the instant plex-breakoff snapshots, refer to the section “[Instant Plex-breakoff Snapshots](#)” on page 168.

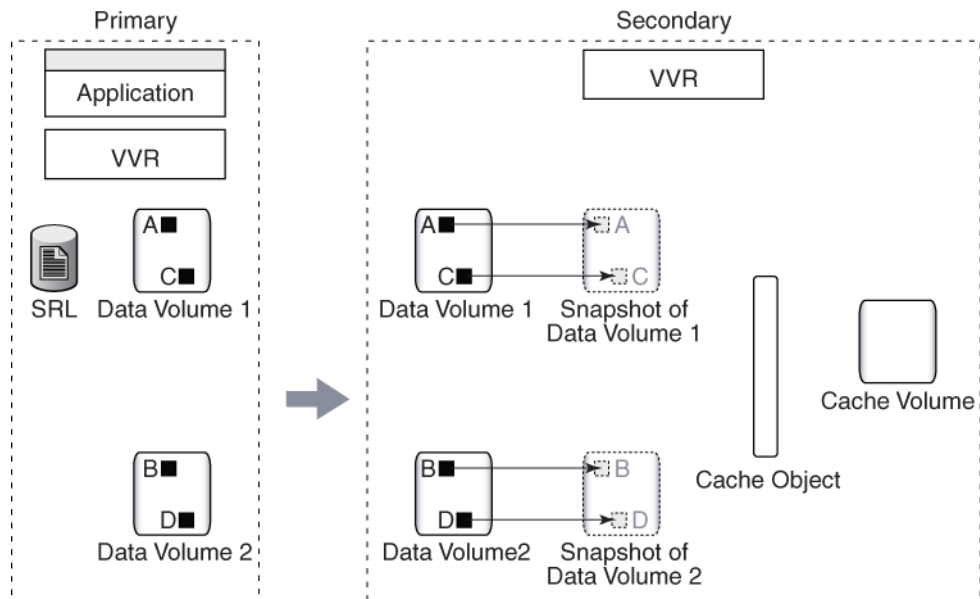
## How VVR Creates Instant Space-Optimized Snapshots

This section explains how VVR creates and manages the space-optimized snapshots.

In the following illustration, the Primary has two data volumes: Data Volume 1 and Data Volume 2. For this example we have indicated two specific blocks, namely, A and C in Data Volume 1 and B and D in Data Volume 2.

The Secondary has two data volumes, Data Volume 1 and Data Volume 2 that have all the replicated data from the Primary, including the specified blocks. The Secondary illustrates an instant space-optimized snapshot in which the data resides on the original volume itself. A Read operation to the snapshots will be redirected to the source volumes and writes will result in a copy-on-write operation. The data will be copied to the snapshots only if there is a write to the original data. Because the snapshots are space-optimized the data will actually get written to the cache object only if there is a write to the original data.

Example 1—How VVR Creates Instant Space-Optimized Snapshots



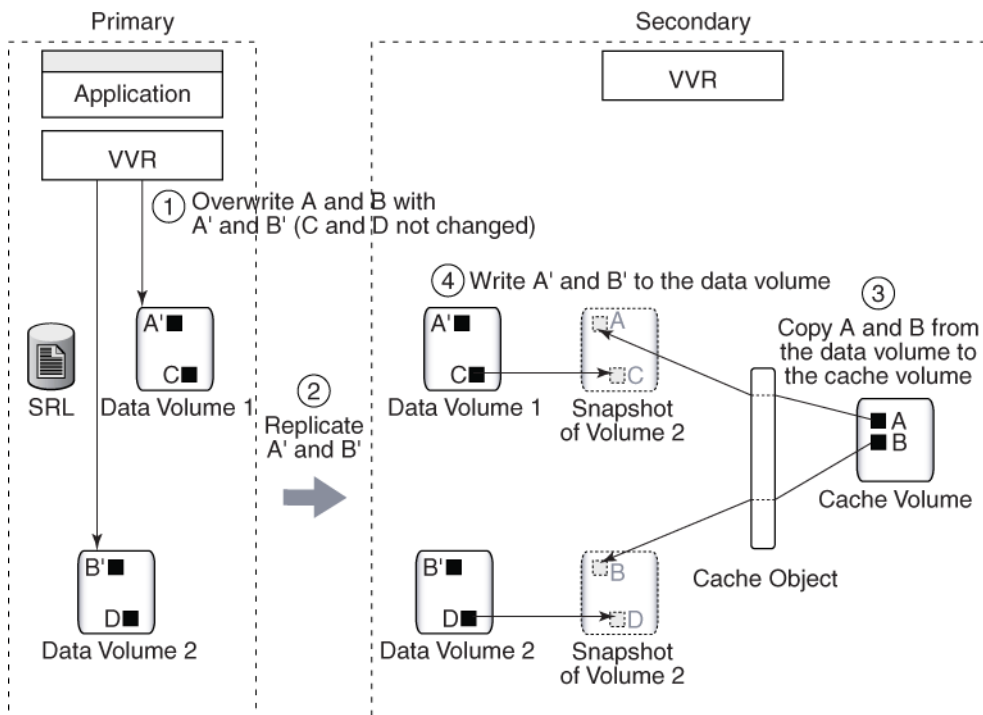
Snap data resides in the original volume



The following illustration indicates the scenario where the Primary receives some updates to the blocks A and B. These are now represented as A' and B'.

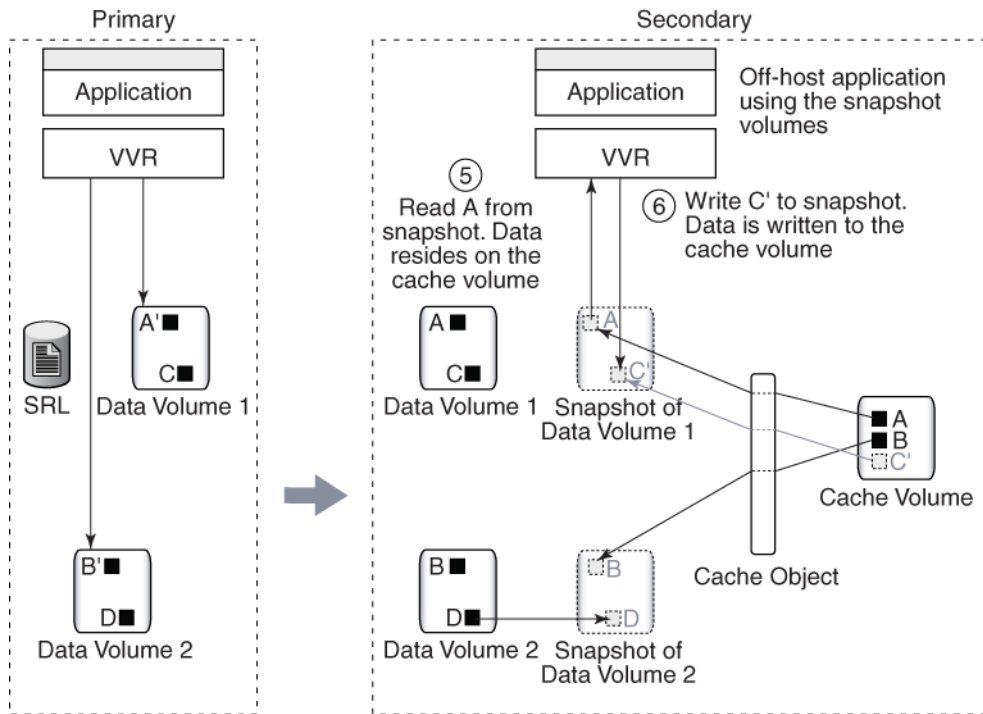
The Secondary also receives the updates A' and B'. The write to the data volumes first results in a copy-on-write on to the space-optimized snapshot. A space-optimized snapshot is created on a cache object that holds all the data for the snapshots. Hence during a copy-on-write, the blocks A and B get written onto the cache object, before the changed blocks are written to the original volumes, Data Volume 1 and Data Volume 2. The cache object stores the blocks persistently in the cache volume after remapping the original offsets.

Example 2—How VVR Writes to the Instant Space-Optimized Snapshots



The following illustration indicates the scenario when there is a write from the application to the block C on the snapshot. This block is now indicated as C'. The changed information for the block C is preserved on the cache object using the copy-on-write mechanism described earlier. If there is a read then the changed block C' will be read from the cache object. Similarly, if there is a request for reading block A from the snapshot volume, it will be fetched from the cache object, where it has been copied earlier.

Example 3—How VVR Reads the Changed Block from the Cache Object



## Comparing the Snapshot Methods Based on Different Features

The following table provides a comparison of the different snapshot methods by feature:

<b>Snapshot Types</b> → <b>Snapshot Features</b> ↓	<b>Traditional Snapshot Feature</b>	<b>Instant Full Snapshot Feature</b>	<b>Instant Space-Optimized Snapshot Feature</b>	<b>Instant Plex Breakoff Snapshot Feature</b>
Requires full storage	Yes	Yes	No	Yes
Requires initialize synchronizing of plexes	Yes	No	No	Yes
Allows snapshots plexes to be reattached (snapback) to the source volume	Yes	Yes	No	Yes
Snapshots can be refreshed	No	Yes	Yes	Yes
Snapshot volumes can be moved into a separate disk group	Yes	Yes	No	Yes
Can be used to restore an RVG to an earlier stable state	No	Yes	Yes	Yes
Can be used as independent volumes	Yes	Yes	No	Yes
Background synchronization	No	Yes	No	No





# Setting Up Replication

# 3

You can configure and administer VERITAS Volume Replicator (VVR) using one of the following interfaces:

- ◆ Command-Line Interface (CLI)

You can use the command-line interface of VVR to configure, administer, and monitor VVR in a distributed environment. The *VERITAS Volume Replicator Administrator's Guide* (this guide) gives instructions on configuring, administering, and monitoring VVR using the CLI.

- ◆ VVR VEA--Java-based desktop GUI

VERITAS Enterprise Administrator (VEA) is a Java-based Graphical User Interface (GUI) that can be used to configure and manage storage objects. VVR VEA enables you to configure, monitor, and administer VVR in a distributed environment. For more information, see "[Administering VVR Using VVR VEA](#)" on page 241.

- ◆ Volume Replicator Web Console (VRW)--Web-based GUI

VERITAS Volume Replicator Web Console (VRW) is the Web-based Graphical User Interface of VVR. For information on configuring, and administering VVR using VRW, see the *VERITAS Volume Replicator Web Console Administrator's Guide*.

This chapter explains how to set up a Replicated Data Set (RDS) using the command-line interface. VVR enables you to set up replication either when the data volumes are zero initialized or contain valid data. Make sure you follow the best practices or recommendations in this chapter to ensure successful configuration of VVR. Read this chapter before you start setting up replication. See "[Examples for Setting Up a Simple Volume Replicator Configuration](#)" on page 73 for detailed examples on how to configure and set up a simple VVR configuration.



Before setting up a Replicated Data Set, lay out your VVR configuration using the “Configuration Worksheet” in the *VERITAS Volume Replicator Installation Guide*. 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.

---

## Best Practices for Setting Up Replication

- ◆ Create one RVG for each application, rather than for each server. For example, if a server is running three separate databases that are being replicated, create three separate RVGs for each database. Creating three separate RVGs helps to avoid write-order dependency between the applications and provides three separate SRLs for maximum performance per application.
- ◆ Create one RVG per disk group. Creating one RVG per disk group enables you to efficiently implement application clustering for high availability, where only one RVG needs to be failed over by the cluster package. If the disk group contains more than one RVG, the applications using the other RVGs would have to be stopped to facilitate the failover. You can use the Disk Group Split feature to migrate application volumes to their own disk groups before associating the volumes to the RVG.
- ◆ Plan the size and layout of the data volumes based on the requirement of your application. You must configure the Primary and Secondary data volumes with the same name.
- ◆ VVR supports VxVM ISP volumes. However, when planning to include these volumes as a part of the RVG make sure that all the volumes in the RVG are of the same type. You can however have RVGs with different types of volumes across hosts. For example, the Primary RVG may have VxVM ISP volumes whereas the secondary RVG may have VxVM non-ISP volumes.
- ◆ Plan the size of the network between the Primary and each Secondary host.
- ◆ Lay out the SRL appropriately to support the performance characteristics needed by the application. Because all writes to the data volumes in an RVG are first written to the SRL, the total write performance of an RVG is bound by the total write performance of the SRL. For example, dedicate separate disks to SRLs and if possible dedicate separate controllers to the SRL.

- ◆ Size the SRL appropriately to avoid overflow. For information on how to determine the size of the SRL, see the *VERITAS Volume Replicator Planning and Tuning Guide*.

---

**Tip** The VERITAS Volume Replicator Advisor (VRAdvisor), a tool to collect and analyze samples of data, can help you determine the optimal size of the SRL. For more information, see the *VERITAS Volume Replicator Advisor User's Guide*.

---

- ◆ Include all the data volumes used by the application in the same RVG. This is mandatory.
- ◆ Provide dedicated bandwidth for VVR over a private network. The RLINK replicates data critical to the survival of the business. Compromising the RLINK compromises the business recovery plan.
- ◆ Use the same names for the data volumes on the Primary and Secondary nodes. If the data volumes on the Primary and Secondary have different names, see [“Mapping the Name of a Secondary Data Volume to a Differently Named Primary Data Volume”](#) on page 401.
- ◆ Make SRLs of equal size and the same names within the same RDS because the Secondary SRL becomes the Primary SRL when the Primary role is transferred.
- ◆ Mirror all data volumes and SRLs. This is optional if you use hardware-based mirroring.
- ◆ The `vradmin` utility creates corresponding RVGs on the Secondary of the same name as the Primary. If you choose to use the `vxmake` command to create RVGs, use the same names for corresponding RVGs on the Primary and Secondary nodes.
- ◆ Associate a DCM to each data volume on the Primary and the Secondary if the DCMs had been removed for some reason. By default, the `vradmin createpri` and `vradmin addsec` commands add DCMs if they do not exist.



## Creating a Replicated Data Set

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

- ✓ Create a Primary Replicated Volume Group (RVG) of the RDS
  - For information on how to associate volume-set component volumes to an RDS, see [“Associating the Component Volumes of a Volume Set to an RDS”](#) on page 127.
- ✓ Add a Secondary to the RDS
- ✓ Change the Replication Settings for the Secondary

The following sections explain each task in detail.

### Creating a Primary RVG of an RDS

The first step in creating an RDS is creating its Primary RVG. VVR enables you to create a Primary RVG of an RDS using the `vradm createpri` command.

The `vradm createpri` command enables you to associate existing data volumes and the Storage Replicator Log (SRL) to the Primary RVG. The `vradm createpri` command performs the following operations:

- ◆ Creates the Primary RVG on the host on which the command is issued.
- ◆ Enables or starts the Primary RVG.
- ◆ Associates DCMs to the data volumes in the RVG.
- ◆ Associates the specified data volumes and SRL to the RVG.

VVR does not support RAID-5 volumes, that is, volumes with usage type `raid5` are not supported. Data volumes must be of usage type `gen` or `fsген`. However, data volumes can be configured on hardware-based RAID-5 disks.

Dirty Region Logs (DRLs) are not needed with VVR because VVR uses the SRL to recover volumes, not the DRLs. If any of the data volumes or the SRL has a DRL, the `vradm createpri` command removes the DRL before the data volume is associated to the RVG.

By default, the `vradm createpri` command adds DCMs to the data volumes, if they have not already been added. The `vradm createpri` command creates the DCM of an appropriate default size based on the size of the volume and mirrors the DCM by default. To create and add a DCM of a size that is different from the default, see [“Associating a Data Change Map to a Data Volume”](#) on page 129. Run the `vradm createpri` command after you have associated the DCMs of the required size to the data volumes.

The `-nodcm` option when used with the `vradm createpri` command associates data volumes to the RVG but does not add DCMs to the data volumes.

If you want to associate additional volumes to the RVG after creating the RVG, use the `vradmin addvol` command. See “[Associating a Volume to a Replicated Data Set](#)” on page 121.

### Prerequisites for creating a Primary RVG of an RDS

- ✓ The data volumes and SRL must exist on the Primary. If the data volumes and SRL do not exist on the Primary, create them. Note that a volume that is to be configured as an SRL cannot be a part of a volume set.
- ✓ The data volumes and SRL must be started. If the data volumes and SRL are not started, start them. When a data volume is started, its state is ACTIVE.
- ✓ The data volumes used by the application must exist in the same RVG. Include the data volumes used by the application in the same RVG.

### ▼ To create a Primary RVG of an RDS

Issue the following command on the host on which you want to create the Primary RVG:

```
# vradmin -g diskgroup createpri rvgname dv01_name,dv02_name... srl_name
```

The argument *rvgname* is the name of the RVG to be created.

The argument *dv01\_name,dv02\_name,...* is a comma-separated list of the names of the data volumes to be associated to the RVG. If the volumes are component volumes of a volume set, then specify the names of the component volumes, not the name of the volume set.

The argument *srl\_name* is the name of the SRL to be associated to the RVG.

Use `-nodcm` option if you do not want DCMs to be added to the data volumes. By default, DCMs are added automatically.

#### Example:

This example shows how to create a Primary RVG `hr_rvg` in the disk group `hrdg`, which contains the data volumes `hr_dv01` and `hr_dv02`, and the volume `hr_srl` that is to be used as the SRL. This example automatically adds DCMs to the data volumes.

```
# vradmin -g hrdg createpri hr_rvg hr_dv01,hr_dv02 hr_srl
```



## Adding a Secondary

After creating the Primary RVG of the RDS, go on to adding a Secondary. Use the `vradm admin addsec` command to add a Secondary RVG to an RDS. This command can also be used to add additional Secondary RVGs. The `vradm admin addsec` command can be issued from any host that is already in the RDS.

---

**Note** If the RDS contains the Primary only, the command must be issued on the Primary. If you issue the `vradm admin addsec` command on the Secondary to be added to the RDS, the command fails as expected.

---

The `vradm admin addsec` command performs the following operations by default:

- ◆ Creates and adds a Secondary RVG of 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. Use the option `-sdg` with the `vradm admin addsec` command to specify a different disk group on the Secondary.
- ◆ If any of the data volumes or the SRL on the Secondary has a DRL, the DRL is removed before the data volume is associated to the RVG. DRLs are not needed with VVR because VVR uses the SRL to recover volumes, not the DRLs.
- ◆ Automatically adds DCMs to the Primary and Secondary data volumes if they do not have DCMs. Use the `-nodcm` option to specify that DCMs are not to be added to the data volumes.

The `vradm admin addsec` command creates the DCM of an appropriate default size based on the size of the volume and mirrors the DCM by default. To create and add a DCM of a size that is different from the default, see [“Associating a Data Change Map to a Data Volume”](#) on page 129. Run the `vradm admin addsec` command after you have associated the DCMs of the required size to the data volumes.

- ◆ 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 to the Primary and Secondary RVGs respectively, the Primary and Secondary RLINKs with default RLINK names `rlk_remotehost_rvgname`. If you choose to use names other than the default, use the `prlink` and `srlink` attributes of the `vradm admin addsec` command to specify the Primary and Secondary RLINK names. For an example of how to specify the RLINK names, see [“Example 2:”](#) on page 49.

### Best practices for adding a Secondary

- ◆ Determine the network and IP addresses to use. Add all participating system names and IP addresses to the `/etc/hosts` files on each system or to the name server database of your name service. Make sure the IP addresses are available (that is, plumbed and up) on the appropriate hosts for your configuration.

- ◆ Plan ahead for application clustering by configuring the IP addresses used for replication as virtual IP addresses. For each replicated data set, the Primary and the Secondary cluster should each have one unique virtual IP address to use as the address for the RLINK. If you do this, you can place VVR under cluster control without having to modify the IP address of the RLINK later. Changing the IP address of an RLINK requires pausing replication.
- ◆ Plan the bandwidth of the network based on your requirement. You can choose to use either the UDP protocol or TCP protocol for network communication between the Primary and Secondary. Also, plan to operate in a firewall environment. For more information, see the *VERITAS Volume Replicator Planning and Tuning Guide*.
- ◆ We recommend that you use the following naming conventions for RLINKs. By default, VVR follows the following naming conventions for RLINKs:
 

Primary RLINK: `rlk_remotehost_rvgrname`. For example: `rlk_london_hr_rvg`

Secondary RLINK: `rlk_remotehost_rvgrname`. For example: `rlk_seattle_hr_rvg`
- ◆ If you plan to have multiple secondaries in your RDS setup, we recommend that you create RLINKs between every pair of secondaries. If you do this, the additional secondaries will be automatically added to the RDS after the migrate operation has completed successfully.
- ◆ Associate a DCM to each data volume on the Primary and the Secondary to use the SRL Protection and Failback Logging features.

### Prerequisites for adding a Secondary

On the Secondary to be added, do the following:

- ✓ Create a disk group with the same name as the Primary disk group.
- ✓ Create data volumes of the same names and lengths as the Primary data volumes.
- ✓ Create an SRL of the same name as the Primary SRL.
- ✓ If any of the volumes in the Primary RVG are component volumes of a volume set, then make sure that the component volumes on the Primary and the Secondary to be added have identical indices.
- ✓ Make sure the `/etc/vx/vras/.rdg` file on the Secondary host to be added to the RDS contains the Primary disk group ID. Ensure that each disk group ID entry in the `.rdg` file is on a separate line. Refer to the `.rdg` file for the sample format for the disk group ID entry.

The `vradm addsec` command checks whether the Primary RVG is authorized to create a corresponding Secondary RVG on the specified Secondary host. A Primary is determined as authorized if the `/etc/vx/vras/.rdg` file on the specified Secondary host contains the Primary disk group ID. If the Primary contains multiple



RVGs in the same disk group, only one entry is required. A plus (+) sign in the `/etc/vx/vras/.rdg` file on the Secondary host indicates that all Primary RVGs on all hosts are authorized to create a Secondary RVG on the specified Secondary host.

The `/etc/vx/vras/.rdg` file on the Secondary host is only used for authorization checking when a Secondary is added, or when remote data volumes are synchronized or verified. To perform these operations after a Secondary takes over from the Primary, the original Primary host should also have an `/etc/vx/vras/.rdg` file containing the disk group ID for the new Primary host.

To display the Primary disk group ID, issue the following command on the Primary host:

```
# vxprint -l diskgroup
```

For example, to enable host `seattle` to create an RVG on Secondary host `london` the `.rdg` file on the host `london` must have the following entries, each on a new line.

```
1083007373.10.seattle
```

#### ▼ To add a Secondary to an RDS

```
# vradmin -g local_diskgroup addsec local_rvgname pri_hostname \  
      sec_hostname
```

The argument `local_diskgroup` is the name of the disk group on the local host.

The argument `local_rvgname` is the name of the RVG on the local host.

The arguments `pri_hostname` and `sec_hostname` are either resolvable hostnames or IP addresses for the Primary and the Secondary hosts. These names are used as `local_host` and `remote_host` attributes while creating RLINKs. The `local_host` and `remote_host` specify the network connection to use for the Primary and Secondary RLINKs.

Use the `-nodcm` option if you do not want to add DCMs to the data volumes. By default, DCMs are automatically added unless the `-nodcm` option is specified.

---

**Note** By default, SRL protection on the new Primary and Secondary RLINKs is set to `autodcm`. If you specify the `-nodcm` option, the `vradmin addsec` command disables SRL protection.

---

Note that the Secondary RVG is added to the disk group with the same name as the Primary disk group, unless specified otherwise using the `-sdg` option.



**Example 1:**

This example shows how to add a Secondary host `london_priv` to the RDS of which `hr_rvg` is a member. For replication, this example uses a private network with the Primary hostname `seattle_priv`, Secondary hostname `london_priv`. On the Secondary, the RVG is added to the same disk group as the Primary, that is, `hrdg`. This example automatically adds DCMs to the data volumes.

```
# vradmin -g hrdg addsec hr_rvg seattle_priv london_priv
```

**Example 2:**

This example shows how to add the Secondary host `london_priv` to the RDS of which `hr_rvg` is a member. It creates the Secondary with the specific Primary and Secondary RLINK names `to_london` and `to_seattle`. The RLINK connects the Primary host `seattle_priv` and the Secondary host `london_priv`. On the Secondary, the RVG is added to the same disk group as the Primary, that is, `hrdg`.

```
# vradmin -g hrdg addsec hr_rvg seattle_priv london_priv \
  prlink=to_london srlink=to_seattle
```

## Changing the Replication Settings for a Secondary

When you add a Secondary to an RDS, the default replication attributes of the Secondary are set to `synchronous=off`, `latencyprot=off`, `srlprot=autodcm`, `packet_size=8400` and `bandwidth_limit=none`. You can set up the replication mode, latency protection, SRL protection, transport protocol, packet size, and the bandwidth used by VVR using the replication attributes, such as `synchronous`, `latencyprot`, and `srlprot`. These attributes are of the form `attribute=value`. Each attribute setting could affect replication and must be set up with care. The `vradmin set` command enables you to change the replication settings between the Primary and a Secondary. This command can be issued from any host in the RDS. It enables you to perform the following tasks:

- ◆ [Setting the Mode of Replication](#)
- ◆ [Setting the Latency Protection](#)
- ◆ [Setting the SRL Overflow Protection](#)
- ◆ [Setting the Network Transport Protocol](#)
- ◆ [Setting the Packet Size](#)
- ◆ [Controlling the Network Bandwidth Used for Replication](#)

The `vradmin set` command changes the corresponding attributes on both the Primary and Secondary RLINK. The three attributes `synchronous`, `latencyprot`, and `srlprot` are only active on the Primary RLINK; however, the Secondary attributes are already set up and ready for use if the Primary role is transferred to the Secondary.



## Setting the Mode of Replication

VVR replicates in two modes: synchronous and asynchronous. The decision to use synchronous or asynchronous mode must be made with an understanding of the effects of this choice on the replication process and the application performance. You can set up VVR to replicate to a Secondary in synchronous or asynchronous mode by setting the `synchronous` attribute of the RLINK to `override`, or `off` respectively.

Setting the `synchronous` attribute to `override` puts the RLINK in synchronous mode. During normal operation, VVR replicates in synchronous mode, but if the RLINK becomes inactive due to a disconnection or administrative action, VVR switches temporarily to asynchronous mode and continues to receive updates from the application and store them in the SRL. After the connection is restored and the SRL is completely drained, the RLINK automatically switches back to synchronous mode. Most system administrators set the `synchronous` attribute to `override`.

The `vradmin` command does not allow you to set the `synchronous` attribute to `fail`. Use the `vxedit` command to set the attribute `synchronous=fail`. For more information on using the `vxedit` command, refer to the `vxedit` manual page.

---

**Caution** You must read the section “Synchronous Mode Considerations” in the *VERITAS Volume Replicator Planning and Tuning Guide* if you use the `synchronous=fail` mode.

---

### ▼ To enable asynchronous mode of replication

To set the replication to asynchronous mode, set the `synchronous` attribute to `off`.

```
# vradmin -g diskgroup set local_rvgname sec_hostname synchronous=off
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host displayed in the output of the `vradmin printrvg` command. If the RDS contains only one Secondary, the argument `sec_hostname` is optional.

**Example:**

To set the mode of replication to asynchronous for the RDS `hr_rvg` between the Primary `seattle` and the Secondary `london`, issue the following command on any host in the RDS:

```
# vradmin -g hrdg set hr_rvg london synchronous=off
```

**▼ To enable synchronous mode of replication**

To set the synchronous attribute of the RLINK to `override`, use the following command:

```
# vradmin -g diskgroup set local_rvgname sec_hostname \  
synchronous=override
```

**Example:**

To set the mode of replication to synchronous for the RDS `hr_rvg` between the Primary `seattle` and the Secondary `london`, issue the following command on any host in the RDS:

```
# vradmin -g hrdg set hr_rvg london synchronous=override
```

**Setting the Latency Protection**

The Latency Protection feature of VVR protects the Secondary host from falling too far behind in updating its copy of data when replicating in asynchronous mode. The Latency Protection feature limits the number of outstanding writes lost in a disaster. It enables automatic control of excessive lag between the Primary and Secondary hosts when replicating in asynchronous mode. For details, see [“Latency Protection—latencyprot attribute”](#) on page 31.

The `vradmin set` command enables you to set the `latencyprot` attribute to `override`, `fail`, or `off`; it also enables you to specify a `latency_high_mark` and a `latency_low_mark`, which indicate when the protection becomes active or inactive.

**▼ To enable latency protection**


---

**Note** Before enabling latency protection, read [“Primary and Secondary Disconnected”](#) on page 32.

---

Set the `latencyprot` attribute to enable latency protection between a Primary and a Secondary:

1. Set the `latencyprot` attribute of the corresponding RLINKs on the Primary and Secondary.



To set the `latencyprot` attribute to override:

```
# vradmin -g diskgroup set local_rvgname sec_hostname \  
latencyprot=override
```

To set the `latencyprot` attribute to fail:

```
# vradmin -g diskgroup set local_rvgname sec_hostname \  
latencyprot=fail
```

2. Set the `latency_high_mark` and the `latency_low_mark` attributes:

```
# vradmin -g diskgroup set local_rvgname sec_hostname \  
latency_high_mark=high_mark
```

```
# vradmin -g diskgroup set local_rvgname sec_hostname \  
latency_low_mark=low_mark
```

The argument `local_rvgname` is the name of the RVG on the local host and represents the RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

Note that the value of `latency_high_mark` must be greater than the value of `latency_low_mark`. We recommend that the difference between the value of `latency_high_mark` and the value of `latency_low_mark` be a small number, for example, 50.

### ▼ To disable latency protection

Setting the `latencyprot` attribute to `off` disables latency protection. This does not limit the number of waiting updates in the SRL.

To set the `latencyprot` attribute to `off`:

```
# vradmin -g diskgroup set local_rvgname sec_hostname latencyprot=off
```

The argument `local_rvgname` is the name of the RVG on the local host and represents the RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

## Setting the SRL Overflow Protection

VVR provides the following modes of SRL overflow protection: `autodcm`, `dcm`, and `override`. For details, see [“Protection Against SRL Overflow—`srlprot` attribute”](#) on page 28. The `vradmin set` command enables you to set up SRL protection by setting the `srlprot` attribute of the corresponding RLINK to either `autodcm`, `dcm`, `override`.

### ▼ To enable SRL overflow protection

To set the `srlprot` attribute to `autodcm`, use the following command:

```
# vradmin -g diskgroup set local_rvgname sec_hostname srlprot=autodcm
```

To set the `srlprot` attribute to `dcm`, use the following command:

```
# vradmin -g diskgroup set local_rvgname sec_hostname srlprot=dcm
```

To set the `srlprot` attribute to `override`, use the following command:

```
# vradmin -g diskgroup set local_rvgname sec_hostname \  
srlprot=override
```

The argument `local_rvgname` is the name of the RVG on the local host and represents the RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.



## Setting the Network Transport Protocol

The value specified for the *protocol* attribute determines the protocol that will be used to communicate between the hosts. You can specify one of the following values for the *protocol* attribute.

- ◆ UDP—The hosts will communicate using UDP/IP protocol. This is the default. If no protocol is specified, then UDP will be used as the protocol of communication between hosts.
- ◆ TCP—The hosts will communicate using TCP/IP protocol.

---

**Note** UDP and TCP are case sensitive.

---

### ▼ To set the network protocol

To set the protocol for RDSs in disk group of version 120 or above, the following `vradmin` command can be used:

```
# vradmin -g diskgroup set local_rvgname sec_hostname \  
    protocol=protocol_name
```

The argument *protocol\_name* is the name of the protocol that the Primary will use to replicate to the Secondary. The protocol can be set to either TCP or UDP.

## Setting the Packet Size

The packet size determines the number of bytes in a packet that are sent to the Secondary host. The packet size can be changed using the *packet\_size* attribute for UDP mode only. If the protocol is set to TCP, the data is sent using the TCP stream. For more information on the *packet\_size* attribute, see the *VERITAS Volume Replicator Planning and Tuning Guide*.

### ▼ To set the packet\_size

```
# vradmin -g diskgroup set local_rvgname sec_hostname packet_size=n
```

The argument *local\_rvgname* is the name of the RVG on the local host and represents the RDS.

The argument *sec\_hostname* is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

The argument *n* represents the packet size in bytes.

The minimum value for the *packet\_size* is 1300 bytes.

The maximum value of the *packet\_size* is 65464 bytes.

**Example:**

To set the packet size between the Primary host `seattle` and the Secondary host `london` to 1400 bytes, issue the following command on any host in the RDS:

```
# vradmin -g hrdg set hr_rvg london packet_size=1400
```

**Controlling the Network Bandwidth Used for Replication**

VVR uses the network to replicate data from the Primary to the Secondary. The Bandwidth Throttling feature enables you to control the maximum network bandwidth to be used by VVR for replication. Bandwidth Throttling controls the rate at which data is sent from the Primary to the Secondary; it does not limit the rate at which the network acknowledgements are sent from the Secondary to the Primary.

You might want to control the bandwidth used by VVR depending on factors such as whether the available network connection is to be used by other applications or exclusively for VVR, the network costs, and network usage over time. For example, if the network is used for purposes other than replication, you might have to control the network bandwidth used by VVR.

Decide on the bandwidth limit for VVR depending on the bandwidth required for VVR, and the bandwidth required for other purposes. For information on how to decide the bandwidth limit to specify for VVR, refer to the *VERITAS Volume Replicator Planning and Tuning Guide*, and the *VERITAS Volume Replicator Advisor User's Guide*.

VVR enables you to change the network bandwidth used for replication to a Secondary even when replication is in progress. It is not necessary to pause replication before you change the bandwidth limit for a Secondary or for an RDS.

Use the `vrstat` command to determine the network bandwidth currently used by VVR. For instructions, see “[Determining the Network Bandwidth Being Used by VVR](#)” on page 98.

Use the `bandwidth_limit` attribute of the `vradmin set` command to set the limit on the network bandwidth used to replicate from the Primary to the Secondary. For example, if `bandwidth_limit` is set to 30 mbps, VVR uses 30 mbps for replication. If `bandwidth_limit` is set to `none`, then VVR uses the available network bandwidth. The default value is `none`.

You can also control the network bandwidth used by VVR when synchronizing volumes, which are not part of an RDS; use the `bandwidth_limit` attribute of the `vradmin syncvol` command to specify the limit.

---

**Note** This value of `bandwidth_limit` specified in the `vradmin syncvol` command is in addition to the bandwidth limit set for replication.

---



For example, if `bandwidth_limit` is set to 30 mbps for replication between a Primary and Secondary in an RDS, and the bandwidth limit to be used when synchronizing volumes that are not part of an RDS using the `vradmin syncvol` command is specified as 10 mbps, then together VVR will use a maximum of 40 mbps.

#### ▼ To control the network bandwidth used for replication

To limit the bandwidth used for replication between the Primary and a Secondary in an RDS, issue the following command on any host in the RDS. In the command, you can either use the units of bandwidth `kbps`, `mbps`, or `gbps`, or abbreviate the units of bandwidth to `k`, `m`, `g`, respectively. The default unit of bandwidth is bits per second (bps).

```
# vradmin -g diskgroup set local_rvgname sec_hostname \  
    bandwidth_limit=value
```

The argument `local_rvgname` is the name of the RVG on the local host and represents the RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

#### Example:

To limit the bandwidth to 30 mbps for the RDS `hr_rvg` between the Primary `seattle` and the Secondary `london`, issue the following command on any host in the RDS:

```
# vradmin -g hrdg set hr_rvg london bandwidth_limit=30mbps
```

#### ▼ To disable Bandwidth Throttling for a Secondary

To disable Bandwidth Throttling for a Secondary in an RDS, issue the following command on any host in the RDS:

```
# vradmin -g diskgroup set local_rvgname sec_hostname \  
    bandwidth_limit=none
```

The argument `local_rvgname` is the name of the RVG on the local host and represents the RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

#### Example:

To disable Bandwidth Throttling for replication between the Primary `seattle` and the Secondary `london` of RDS `hr_rvg`, issue the following command on any host in the RDS:

```
# vradmin -g hrdg set hr_rvg london bandwidth_limit=none
```



### ▼ To control the network bandwidth used to synchronize volumes

To limit the network bandwidth used by VVR when synchronizing volumes that are not part of an RDS, issue the following command:

```
# vradmin -g diskgroup syncvol local_vols_list \  
    remote_hostname...bandwidth_limit=value
```

The argument *local\_vols\_list* is a comma-separated list of volumes on the local host. The names of the volumes on the local and remote hosts are assumed to be the same.

The argument *remote\_hostname* is a space-separated list of names of the remote hosts on which the volumes to be resynchronized reside. It must be possible for IP to resolve the remote host names.

#### Example:

This example shows how to limit the network bandwidth used by VVR when using full synchronization to synchronize the remote volumes on host `london` with the local volumes `hr_dv01`, `hr_dv02`, `hr_dv03` in the disk group `hrdg` on the local host `seattle`. The names of the disk group and the volumes on the remote host are the same as the names of the disk group and volumes on the local host.

```
# vradmin -g hrdg -full syncvol hr_dv01,hr_dv02,hr_dv03 london /  
    bandwidth_limit=10mbps
```



## Synchronizing the Secondary and Starting Replication

This section explains how to synchronize the Secondary and start replication.

### Methods to Synchronize the Secondary

You can synchronize the Secondary using the network, using block-level tape backup or by physically moving disks to the Secondary. Use one of the following methods to synchronize the Secondary depending on your environment:

- ◆ Using the network
  - ◆ Automatic synchronization
  - ◆ Full synchronization with checkpoint
  - ◆ Difference-based synchronization with checkpoint
- ◆ Using block-level tape backup
  - ◆ Block-level tape backup and checkpointing
- ◆ Moving disks physically
  - ◆ Disk Group Split and Join

The following tables explain when and how to use the different synchronization methods:

## Using the Network

You can synchronize the Secondary over the network either when the application is active or inactive.

To Synchronize the Secondary:	Perform:	Using This Command:
completely	automatic synchronization and start replication	<code>vradmin -a startrep</code>
completely	full synchronization with checkpoint	<code>vradmin -full -c <i>checkpoint</i> syncrvg</code>
when there is little difference between the data on the Primary and Secondary data volumes of the RDS	difference-based synchronization with checkpoint. For more information, see <a href="#">“Using Difference-Based Synchronization”</a> on page 70.	<code>vradmin -c <i>checkpoint</i> syncrvg</code>



## Using Block-level Tape Backup

---

<b>To Synchronize the Secondary:</b>	<b>Do the following:</b>	<b>Using This Command:</b>
completely and when a large amount of data must be moved from the Primary to the Secondary	<ol style="list-style-type: none"><li data-bbox="568 388 815 446">1. Start a Primary checkpoint.</li><li data-bbox="568 539 815 597">2. Perform a block-level backup of the Primary.</li><li data-bbox="568 630 815 687">3. End the Primary checkpoint.</li><li data-bbox="568 720 815 868">4. Restore the tapes on the Secondary and start replication to the Secondary using the checkpoint.</li></ol>	<pre data-bbox="829 388 1270 418">vxrvg -c <i>checkpoint</i> checkstart</pre> <pre data-bbox="829 630 1270 659">vxrvg checkend</pre> <pre data-bbox="829 720 1270 748">vradmin -c <i>checkpoint</i> startrep</pre>

---



## Moving Disks Physically

To Synchronize the Secondary:	Use This Feature	Using This Command:
completely by physically moving disks from the location of the Primary host to the location of Secondary host	Disk Group Split and Join	For instructions, see <a href="#">“Using the Disk Group Split and Join Feature”</a> on page 68.

## Using the Automatic Synchronization Feature

The Automatic Synchronization feature enables you to transfer the data on the Primary to the Secondary over the network. You can synchronize the Secondary using automatic synchronization either when the application is active or inactive.

The Automatic Synchronization procedure transfers data in the Primary data volumes to the Secondary by reading the Primary data volumes from start to finish and sending the data to the Secondary.

**Note** Automatic Synchronization does not maintain the order of writes; therefore, the Secondary is inconsistent until the process is complete.

The Secondary becomes consistent after the automatic synchronization completes. To use Automatic Synchronization successfully, the network must be sized appropriately. Note that the synchronization will complete only if the Primary receives writes at a lesser rate than they can be sent to the Secondary. If the Primary receives writes at a faster rate than they can be sent to the Secondary, the synchronization might never complete, especially if the writes are dispersed widely in the volume.

This feature enables you to synchronize multiple Secondary hosts at the same time. When performing automatic synchronization to multiple Secondary hosts, synchronization proceeds at the rate of the slowest network.

VVR pauses synchronization if the Secondary fails or the network disconnects. If the Primary fails while synchronization is in progress, the synchronization continues from the point at which it had stopped when the Primary recovers.



### Prerequisite for using Automatic Synchronization

- ◆ Each data volume in the Primary RVG must have a DCM associated to it. If data volumes do not have DCMs, an attempt to automatically synchronize a Secondary fails.

The `vradmin startrep` command when used with the option `-a` enables you to start replication and automatically synchronize the Secondary data volumes with the Primary data volumes in an RDS; it brings the Secondary data volumes up-to-date with the Primary data volumes. You can use this command 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.

The `vradmin startrep` command can be issued from any host in the RDS. To check the status and progress of the automatic synchronization, use the `vxrlink status` command on the Primary RLINK. For more information, see [“Displaying the Status of a Secondary”](#) on page 89.

To synchronize the Secondary and start replication using automatic synchronization, issue the following command:

```
# vradmin -g diskgroup -a startrep local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host displayed in the output of the `vradmin printrvg` command. If the RDS contains only one Secondary, the `sec_hostname` is optional.

### Example—Using the Automatic Synchronization Feature

In this example, the data volumes in the Primary RVG `hr_rvg` on host `seattle` contain valid data and the application is active. To start replication and synchronize the Secondary RVG `hr_rvg` on host `london`, issue the following command:

```
# vradmin -g hrdg -a startrep hr_rvg london
```

**Notes on Using Automatic Synchronization**

- ◆ If you associate a new volume to an RDS while automatic synchronization is in progress, VVR does not automatically synchronize the newly associated data volume.
- ◆ In an RDS containing multiple Secondaries that have SRL overflow protection set to dcm, more than one Secondary may require the use of the DCM. If one Secondary is undergoing automatic synchronization and the RLINK of another Secondary is about to overflow, the Automatic Synchronization is abandoned and the DCM becomes active for the overflowing RLINK.
- ◆ If you try to automatically synchronize a new RLINK while an existing RLINK is using the DCM mechanism, the automatic synchronization fails.
- ◆ To remove a Secondary from a DCM resynchronization process, detach the corresponding Primary RLINK.
- ◆ If you try to dissociate a DCM from a data volume while the DCM is in use, the operation fails.
- ◆ If the DCM is detached because of I/O errors while the DCM is in use, the resynchronization is abandoned and the RLINKs that are being synchronized are detached.



## Using the Full Synchronization Feature

This section explains how to use the Full Synchronization feature of VVR to synchronize the Secondary completely and start replication. Full synchronization compresses zeroes while processing the data and hence proves beneficial when a large part of the Primary data volumes contain zeroes. However, we recommend that you use Automatic Synchronization to synchronize the Secondary because its performance is better than Full Synchronization. Automatic Synchronization also handles network outages efficiently and continues even after the system reboots.

Full synchronization synchronizes the Secondary over the network when the Primary data volumes contain data and when the application is active or inactive. After the Primary and Secondary are synchronized, replication must be started.

By default, the `vradmin syncrvg` command synchronizes Secondary data volumes using difference-based synchronization. To perform a full synchronization, specify the `-full` option.

We recommend that you always use the `-c` option with the `vradmin syncrvg` command to synchronize the Secondary using full synchronization. The `-c checkpoint` option starts a checkpoint, synchronizes the data volumes, and ends the checkpoint after the synchronization completes. After the `vradmin syncrvg` command completes, use this checkpoint with the `vradmin startrep` command to start replication. To delete the Primary checkpoints, use the `vrxvg checkdelete` command.

The SRL must be large enough to hold the incoming updates to the Primary data volumes while the synchronization is in progress. The SRL might fill up and the checkpoint might overflow if the number of writes to the Primary data volumes is high when the Secondary is being synchronized.

A checkpoint that has overflowed becomes invalid and cannot be used to start replication. If the checkpoint overflows while synchronization is in progress, the `vradmin syncrvg` command must be issued again.

The `vradmin syncrvg` command can be used to synchronize multiple Secondaries at the same time. The `vradmin syncrvg` command displays the progress of the synchronization.



▼ **To synchronize the Secondary RVG with the Primary RVG using full synchronization with checkpoint**

---

**Note** The `vradmin syncrvg` command cannot be used for an RVG that has volume-set component volumes associated to it.

---

1. Verify that the RLINKs are detached to ensure that replication is stopped.
2. To synchronize the Secondary RVG, issue the following command:

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

Note that you can use the `-c` option with the `vradmin syncrvg` command when performing full synchronization, to automatically start a checkpoint with the specified name. After the data volumes are synchronized the checkpoint is ended. This checkpoint can then be used to start replication using the `vradmin startrep` command.

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname...` is a space-separated list of the names of the Secondary hosts as displayed in the output of the `vradmin printrvg` command.

The argument `checkpoint_name` specifies the name of the Primary checkpoint of your choice.

3. After the synchronization completes, start replication to the Secondary with the checkpoint:

```
# vradmin -g diskgroup -c checkpoint_name startrep local_rvgname \  
    sec_hostname
```

After the RLINKs are attached, the Secondary remains inconsistent until it has received all of the accumulated updates up to the checkend. While the Secondary is inconsistent, the `inconsistent` flag is set on the Secondary RLINK. After all updates up to the checkend have been received and applied at the Secondary, the `inconsistent` flag is cleared.

To view the status of synchronization, see [“Displaying the Status of a Secondary”](#) on page 89.



### Example—Synchronizing the Secondary Using Full Synchronization with Checkpoint

This example explains how to synchronize the Secondary RVG `hr_rvg` on the Secondary host `london` with the Primary RVG on host `seattle`.

▼ **To synchronize the Secondary RVG `hr_rvg` on `london` with its Primary RVG on `seattle` using full synchronization**

1. Verify that the RLINKs are detached to ensure that replication is stopped.
2. Issue the following command from any host in the RDS:

```
# vradmin -g hrdg -full -c checkpt_presync syncrvg hr_rvg london
```

Note that you can use the `-c` option with the `vradmin syncrvg` command when performing full synchronization, to automatically start a checkpoint with the specified name. After the data volumes are synchronized the checkpoint is ended. This checkpoint can then be used to start replication using the `vradmin startrep` command.

The name `checkpt_presync` is the Primary checkpoint that you will create.

3. After the synchronization is complete, issue the following command to start replication using the checkpoint:

```
# vradmin -g hrdg -c checkpt_presync startrep hr_rvg london
```

## Using Block-level Backup and Checkpoint

This method is useful for low bandwidth networks or very large data sets. You can use the block-level backup and checkpoint method to synchronize the Secondary when a backup of the data is available and a checkpoint has been started on the Primary. You do not have to use the network to transfer the data. This method does have a risk of SRL overflow.

Make sure that the SRL is large enough to contain all the writes made by the application while synchronization is in progress. For instructions on how to resize the SRL, see [“Resizing the SRL”](#) on page 141.

---

**Caution** During the process the checkpoint will overflow if the SRL fills up. To determine if the checkpoint has overflowed, issue the `vrvrg cplist rvg_name` command on the Primary to display the list of valid checkpoints.

---

For detailed instructions, see [“Example—Synchronizing the Secondary Using Block-level Backup”](#) on page 67.

### ▼ To synchronize the Secondary using backup and Primary checkpoint

1. Start a Primary checkpoint using the `vrxvg checkstart` command:

```
# vxrvg -g diskgroup -c checkpoint_name checkstart local_rvgname
```

2. Perform a block-level backup of the data volumes in the Primary RVG.

3. End the checkpoint in the SRL when the backup is complete by using the `vrxvg checkend` command:

```
# vxrvg -g diskgroup checkend local_rvgname
```

4. Restore the backup to the Secondary data volumes.

5. Start replication using the checkpoint after the restore on the Secondary completes:

```
# vradmin -g diskgroup -c checkpoint_name startrep local_rvgname \
  sec_hostname
```

After the RLINKs are attached, the Secondary remains inconsistent until it has received all of the accumulated updates up to the `checkend`. While the Secondary is inconsistent, the `inconsistent` flag is set on the Secondary RLINK. After all updates up to the `checkend` have been received and applied at the Secondary, the `inconsistent` flag is cleared.

### Example—Synchronizing the Secondary Using Block-level Backup

This example explains how to synchronize the Secondary RVG `hr_rvg` on the Secondary host `london` with the Primary RVG on host `seattle` using block-level backup and checkpoint.

1. Start a Primary checkpoint on `seattle`:

```
# vxrvg -g hrdg -c checkpoint_presync checkstart hr_rvg
```

2. Perform a block-level backup of the data volumes in the Primary RVG.

3. End the Primary checkpoint when the backup is complete:

```
# vxrvg -g hrdg checkend hr_rvg
```

4. Restore the backup to the Secondary data volumes.

5. Start replication using the checkpoint after the restore is complete:

```
# vradmin -g hrdg -c checkpoint_presync startrep hr_rvg london
```



## Using the Disk Group Split and Join Feature

The Disk Group Split and Join feature of VERITAS Volume Manager enables you to synchronize the Secondary. For more information on the Disk Group Split and Join feature, refer to the *VERITAS Volume Manager Administrator's Guide*. To set up replication using this method, ensure that you have a valid Disk Group Split and Join license on your system. For detailed instructions, see “[Example 4—Setting Up Replication Using Disk Group Split and Join](#)” on page 78.

### ▼ To synchronize the Secondary using Disk Group Split and Join

Perform the following steps on the Primary:

1. Create a snapshot plex for each data volume in the Primary RVG by issuing the following command on the Primary:

```
# vxassist -g diskgroup snapstart dv_name
```

You can use the `-b` option with the `vxassist snapstart` command to run the command in the background. Note that if you use the `-b` option of the `vxassist snapstart` command, you must wait for the snapshot plexes for all the data volumes in the RVG to be created and synchronized completely before you proceed to the next step. When the plex synchronization completes, the output of the `vxprint` command displays the state of the new snapshot plex as `SNAPDONE`.

2. Start a Primary checkpoint by issuing the following command on the Primary:

```
# vxrvg -g diskgroup -c checkpoint_name checkstart local_rvgname
```

3. Take a snapshot of each data volume in the Primary RVG by issuing the following command on the Primary:

```
# vxrvg -g diskgroup snapshot local_rvgname
```

4. End the checkpoint by issuing the following command on the Primary:

```
# vxrvg -g diskgroup checkend local_rvgname
```

5. Split the snapshot volumes into a new disk group by issuing the following command on the Primary:

```
# vxdg split diskgroup new_diskgroup SNAP-dv_name ...
```

6. Rename each snapshot volume in the new disk group with the same name as the corresponding data volume in the Primary RVG by issuing the following command on the Primary:

```
# vxedit -g new_diskgroup rename SNAP-dv_name dv_name
```

7. Deport the split-off disk group, rename it to the same name as the disk group of the Primary RVG, and change the ownership of the split-off disk group to the Secondary host so that it may be automatically imported on the Secondary on reboot.

```
# vxvg -n diskgroup -h sec_hostname deport new_diskgroup
```

The argument *sec\_hostname* is the name of the Secondary host displayed in the output of the `uname` command.

8. Physically remove the disks contained in the deported disk group by following the procedures recommended by the disk manufacturer; then attach the disks to the Secondary host.
9. On the Secondary, import the disks that were moved over from the Primary if not already imported:

```
# vxvg import diskgroup
```

10. Add the Secondary to the RDS by issuing the following command on the Primary:

```
# vradm -g diskgroup addsec local_rvgname pri_hostname \
    sec_hostname
```

11. Start replication by issuing the following command from any host in the RDS:

```
# vradm -g diskgroup -c ckpt_name startrep local_rvgname \
    sec_hostname
```

The argument *sec\_hostname* is the name of the Secondary host displayed in the output of the `vradm printrvg` command. If the RDS contains only one Secondary, the *sec\_hostname* is optional.



## Using Difference-Based Synchronization

You can synchronize the Secondary using difference-based synchronization when there is little difference between the Primary and Secondary data volumes in an RDS.

Difference-based synchronization can be used to transfer data over the network when the application is active or inactive.

In difference-based synchronization, the `syncrvg` command generates MD5 checksums for the data blocks on the Primary data volume and the corresponding Secondary data volume and compares these checksums. The `syncrvg` command then transfers over the network only those blocks for which checksums do not match. These steps are repeated for the entire Primary data volume and Secondary data volume.

MD5 checksum is generated using the RSA Data Security, Inc. MD5 Message-Digest Algorithm.

Difference-based synchronization is useful in situations such as:

- ◆ Storage Replicator Log (SRL) Overflow— To synchronize the Secondary when the SRL protection has not been set to use the Data Change Map (DCM).
- ◆ Failing Back to the Original Primary—To synchronize the original Primary data volumes with the new Primary data volumes.

The `vradmin syncrvg` command enables you to synchronize the Secondary RVG with the Primary RVG based on differences. You can issue the `vradmin syncrvg` command from any host in the RDS. The `vradmin syncrvg` command synchronizes the data volumes associated with the Secondary RVG in an RDS with the corresponding data volumes associated with its Primary RVG. The `vradmin syncrvg` command can be used to synchronize multiple Secondaries at the same time.

### ▼ To synchronize Secondary RVG with Primary RVG based on differences

Before issuing this command, verify that the RLINKs are detached. Use the `-c` checkpoint option with the `vradmin syncrvg` command as follows:

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

Note that you can use the `-c` option with the `vradmin syncrvg` command when performing difference-based synchronization, to automatically start a checkpoint with the specified name. After the data volumes are synchronized the checkpoint is ended. This checkpoint can then be used to start replication using the `vradmin startrep` command.

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is a space-separated list of the names of the Secondary hosts as displayed in the output of the `vradmin printrvg` command.

The argument *ckpt\_name* specifies the name of the Primary checkpoint of your choice.

### Example—Synchronizing the Secondary Based on Differences

This example explains how to synchronize the Secondary RVG *hr\_rvg* on the Secondary host *london* with the Primary RVG on host *seattle*.

#### ▼ To synchronize the Secondary RVG, *hr\_rvg*, on *london* with its Primary RVG on *seattle* based on differences

Before issuing this command, make sure that the RLINKs are detached.

```
# vradmin -g hrdg -c ckpt_presync syncrvg hr_rvg london
```

Note that you can use the *-c* option with the `vradmin syncrvg` command when performing difference-based synchronization to automatically start a checkpoint with the specified name. After the data volumes are synchronized the checkpoint is ended. This checkpoint can then be used to start replication using the `vradmin startrep` command.

The name *ckpt\_presync* is the Primary checkpoint that you will create.



## Starting Replication When the Data Volumes are Zero Initialized

Use the option `-f` with the `vradmin startrep` command to start replication when the Primary and Secondary data volumes are zero initialized. The `vradmin startrep` command can be issued from any host in an RDS.

To start replication to a Secondary in an RDS when the data volumes are zero initialized:

```
# vradmin -g diskgroup -f startrep local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command. If the RDS contains only one Secondary, the `sec_hostname` is optional.

### Example:

To start replication from the Primary RVG `hr_rvg` on `seattle` to the Secondary RVG on host `london` when the data volumes are zero initialized, issue the following command from any host in the RDS:

```
# vradmin -g hrdg -f startrep hr_rvg london
```



## Examples for Setting Up a Simple Volume Replicator Configuration

The examples in this section explain how to use VERITAS Volume Replicator (VVR) to set up a simple VVR configuration under different situations. The examples explain how to set up a VVR configuration with one Secondary and hence one RLINK; however, VVR enables you to configure and set up configurations with multiple Secondaries. The examples give the steps to replicate from the Primary host `seattle` to the Secondary host `london`.

### Assumptions:

- ◆ These examples assume that the Primary `seattle` and Secondary `london` have a disk group named `hrdg` with enough free space to create the VVR objects mentioned in the example.
- ◆ Examples 1, 2, 3, 4, and 5 assume that the Primary data volumes have been set up and contain data.
- ◆ The examples assume that the `/etc/vx/vras/.rdg` file on the Secondary host contains the Primary diskgroup ID. Ensure that each disk group ID entry in the `.rdg` file appears on a separate line. A Secondary can be added to an RDS only if the `/etc/vx/vras/.rdg` file on the Secondary host contains the Primary disk group ID. Use the `vxprint -l diskgroup` command to display the disk group ID of the disk group `hrdg`, which is being used.

### Configuration Considerations:

Consider the following in each example:

- ✓ The data volumes on the Secondary must have the same names and sizes as the data volumes on the Primary.
- ✓ The name of the Storage Replicator Log (SRL) on the Secondary must be the same as the name of the SRL on the Primary.
- ✓ The SRL must be created on disks that do not have other volumes on them.
- ✓ The data volumes and SRL must be mirrored.

In the examples, each data volume is 4 GB; the Primary and Secondary SRL are 4 GB each.



The examples in this chapter use the following names:

Primary Hostname: `seattle`

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

Secondary Hostname: `london`

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

## Creating a Replicated Data Set for the Examples

Perform the following steps to create a Replicated Data Set (RDS) for all the examples in this section, except “[Example 4—Setting Up Replication Using Disk Group Split and Join](#)” on page 78.

---

**Note** This example assumes that the disks available on the system have labels such as disk01, disk02, disk03 and so on.

---

1. Create the data volumes on the Secondary host `london`. Use different disks for the data volumes and SRL.

```
# vxassist -g hrdg make hr_dv01 4G \  
    layout=mirror logtype=dcm mirror=2 disk01 disk02  
# vxassist -g hrdg make hr_dv02 4G \  
    layout=mirror logtype=dcm mirror=2 disk03 disk04
```

2. Create the SRL on disks that do not have other volumes on them by typing the following command on the Primary `seattle` and the Secondary `london`:

```
# vxassist -g hrdg make hr_srl 4G mirror=2 disk05 disk06
```

---

**Note** You must create the SRL on disks that do not contain any other volume.

---

3. Create the Primary RVG of the RDS by typing the following command on the Primary `seattle`:

```
# vradm -g hrdg createpri hr_rvg hr_dv01,hr_dv02 hr_srl
```

4. Make sure the `/etc/vx/vras/.rdg` file on the Secondary host `london` contains the Primary disk group ID of `hrdg`; then add the Secondary `london` to the RDS by typing the following command on the Primary `seattle`:

```
# vradm -g hrdg addsec hr_rvg seattle london
```



## Example 1—Setting Up Replication Using Automatic Synchronization

This example assumes that the RDS has been created using the procedure “[Creating a Replicated Data Set for the Examples](#)” on page 75. You can synchronize the Secondary using automatic synchronization when the application is active or inactive.

- ❖ Start Secondary synchronization and replication using automatic synchronization by issuing the following command from any host in the RDS:

```
# vradmin -g hrdg -a startrep hr_rvg london
```

## Example 2—Setting Up Replication Using Full Synchronization

This example assumes that the RDS has been created using the procedure “[Creating a Replicated Data Set for the Examples](#)” on page 75. You can synchronize the Secondary using full synchronization with checkpoint when the application is active or inactive.

1. Synchronize the Secondary RVG `hr_rvg` on `london` with its Primary RVG on `seattle` using full synchronization with checkpoint:

```
# vradmin -g hrdg -full -c chkpt_presync syncrvg hr_rvg london
```

2. Start replication with checkpoint by issuing the following command on any host in the RDS:

```
# vradmin -g hrdg -c chkpt_presync startrep hr_rvg london
```

## Example 3—Setting Up Replication Using Block-level Backup and Checkpointing

This example assumes that the RDS has been created using the procedure “[Creating a Replicated Data Set for the Examples](#)” on page 75. You can synchronize the Secondary using block-level backup and checkpointing when the application is active or inactive.

1. Start a checkpoint on the Primary:

```
# vxrvrg -g hrdg -c checkpt_presync checkstart hr_rvg
```

Note down the checkpoint name you use, that is, `checkpt_presync`.

2. Perform a block-level backup of the data volumes in the Primary RVG.
3. End the Primary checkpoint when the backup is complete:

```
# vxrvrg -g hrdg checkend hr_rvg
```

4. Restore the backup to the Secondary data volumes.
5. Use the `vxrvrg cplist` command on the Primary to check whether the checkpoint you created is still valid. If the checkpoint has overflowed, repeat [step 1](#) to [step 4](#).

The output resembles:

Name	MBytes	% Log	Started/Completed
checkpt_presync	10	9	Completed

6. Start replication using the checkpoint:

```
# vradmin -g hrdg -c checkpt_presync startrep hr_rvg london
```

7. On the Primary, check whether the `consistent` flag is set on the Primary RLINK using the `vxprint` command. The RLINK becomes `consistent` only after the data contained in the checkpoint is sent to the Secondary. Wait and then issue the following command on the Primary:

```
# vxprint -g hrdg -l rlk_london_hr_rvg
```

If the Secondary is consistent, the synchronization was successful.

If the checkpoint overflows before the Secondary becomes consistent, the synchronization process has failed. Increase the size of the SRL, and then restart the procedure beginning at [step 1](#). To resize the SRL, see “[Resizing the SRL](#)” on page 141.



It is likely that there might be writes beyond the checkpoint that are yet to be sent to the Secondary after `consistent` flag is set on the RLINK. Use the `vxrlink status` command to check whether the RLINK is up-to-date:

```
# vxrlink -g hrdg status rlk_london_hr_rvg
```

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. Eventually, any checkpoint that becomes STALE is unusable. There is no warning to indicate that this has occurred. However, the `vxrvg cplist` command indicates that the checkpoint has overflowed and hence is unusable. For more information, see the section “[Displaying a List of Checkpoints](#)” on page 91.

## Example 4—Setting Up Replication Using Disk Group Split and Join

Do not perform the steps listed in “[Creating a Replicated Data Set for the Examples](#)” on page 75.

1. Create the SRL on disks that do not have other volumes on them by typing the following command on the Primary `seattle`:

```
# vxassist -g hrdg make hr_srl 4G mirror=2
```

---

**Note** You must create the SRL on disks that do not contain any other volume.

---

2. Create the Primary RVG of the RDS by typing the following command on the Primary `seattle`:

```
# vradmin -g hrdg createpri hr_rvg hr_dv01,hr_dv02 hr_srl
```

3. Create a snapshot plex for each data volume in the Primary RVG by issuing the following commands on the Primary `seattle`:

```
# vxassist -g hrdg snapstart hr_dv01
```

```
# vxassist -g hrdg snapstart hr_dv02
```

You can use the `-b` option with the `vxassist snapstart` command to run the command in the background. Note that if you use the `-b` option of the `vxassist snapstart` command, you must wait for the snapshot plexes for all the data volumes in the RVG to be created and synchronized completely before you proceed to the next step. When the plex synchronization completes, the output of the `vxprint` command displays the state of the new snapshot plex as `SNAPDONE`.

4. Start a Primary checkpoint by issuing the following command on the Primary `seattle`:  

```
# vxrvg -g hrdg -c checkpt_presync checkstart hr_rvg
```
5. Take snapshot of each data volume in the Primary RVG by issuing the following command on the Primary `seattle`:  

```
# vxrvg -g hrdg snapshot hr_rvg
```
6. End the checkpoint by issuing the following command on the Primary `seattle`:  

```
# vxrvg -g hrdg checkend hr_rvg
```
7. Split the snapshot volumes into a new disk group by issuing the following command on the Primary `seattle`:  

```
# vxdg split hrdg new_hrdg SNAP-hr_dv01 SNAP-hr_dv02
```
8. Rename each snapshot volume in the new disk group with the same name as the corresponding data volume in the Primary RVG by issuing the following commands on the Primary `seattle`:  

```
# vxedit -g new_hrdg rename SNAP-hr_dv01 hr_dv01  
# vxedit -g new_hrdg rename SNAP-hr_dv02 hr_dv02
```
9. Deport the split-off disk group, rename it to the same name as the disk group of the Primary RVG and change the ownership of the split-off disk group to be the Secondary host so that it may be automatically imported on the Secondary on reboot. To do this, issue the following command on the Primary `seattle`:  

```
# vxdg -n hrdg -h london deport new_hrdg
```
10. Physically remove the disks contained in the deported disk group by following the procedures recommended by the disk manufacturer; then attach the disks to the Secondary host.
11. On the Secondary `london`, import the disks that were moved over from the Primary if not already imported.  

```
# vxdg import hrdg
```
12. Create the SRL on disks that do not have other volumes on them by typing the following command on the Secondary `london`:  

```
# vxassist -g hrdg make hr_srl 4G mirror=2
```

---

**Note** You must create the SRL on disks that do not contain any other volume.

---



13. Make sure the `/etc/vx/vras/.rdg` file on the Secondary host `london` contains the Primary disk group ID of `hrdg`; then add the Secondary to the RDS by issuing the following command on the Primary `seattle`:

```
# vradmin -g hrdg addsec hr_rvg seattle london
```

14. Start replication by issuing the following command on any host in the RDS:

```
# vradmin -g hrdg -c checkpt_presync startrep hr_rvg london
```

## Example 5—Setting Up Replication Using Differences-Based Synchronization

This example assumes that the RDS has been created using the procedure “[Creating a Replicated Data Set for the Examples](#)” on page 75. You can synchronize the Secondary using difference-based synchronization with checkpoint when the application is active or inactive.

1. Synchronize the Secondary RVG `hr_rvg` on `london` with its Primary RVG on `seattle` using difference-based synchronization with checkpoint:

```
# vradmin -g hrdg -c chkpt_presync syncrvg hr_rvg london
```

2. Start replication with checkpoint by issuing the following command on any host in the RDS:

```
# vradmin -g hrdg -c chkpt_presync startrep hr_rvg london
```



## Example 6—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.

1. Create the data volumes by typing the following commands on the Primary `seattle` and Secondary `london`. Use different disks for the data volumes and SRL.

```
# vxassist -g hrdg make hr_dv01 4G \  
  layout=mirror logtype=dcm mirror=2 init=zero disk01 disk02  
  
# vxassist -g hrdg make hr_dv02 4G \  
  layout=mirror logtype=dcm mirror=2 init=zero disk03 disk04
```

2. Create the SRL on disks that do not have other volumes on them by typing the following command on the Primary `seattle` and Secondary `london`:

```
# vxassist -g hrdg make hr_srl 4G mirror=2 disk05 disk06
```

3. Create the Primary RVG of the RDS by typing the following command on the Primary `seattle`:

```
# vradmin -g hrdg createpri hr_rvg hr_dv01,hr_dv02 hr_srl
```

4. Make sure the `/etc/vx/vras/.rdg` file on the Secondary host `london` contains the Primary disk group ID of `hrdg`; then add the Secondary `london` to the RDS by typing the following command on the Primary `seattle`:

```
# vradmin -g hrdg addsec hr_rvg seattle london
```

---

**Note** Do not start the application or mount the file system before completing [step 5](#).

---

5. Start replication using the option `-f` with the `vradmin startrep` command by typing the following command on any host in the RDS:

```
# vradmin -g hrdg -f startrep hr_rvg london
```

After completing this step, start the application.





# Displaying Configuration Information

---

# 4

This chapter describes the commands to display a consolidated view of the Replicated Data Sets (RDSs) in a VERITAS Volume Replicator (VVR) environment. These display commands can be entered from any host in an RDS.

The `vradmin print` commands display the corresponding objects on all hosts. The `vxprint` print commands display detailed information about a specific object on the host on which the command is issued.

## Using the vradmin Display Commands

This section describes the `vradmin print` commands.

### Displaying RVG and RDS Information

Use the `vradmin printrvg` command to display information about the RDSs on a host. You can run the `vradmin printrvg` command from any host in an RDS.

To display information about a specific RDS, on a local host type:

```
# vradmin -g diskgroup printrvg local_rvgname
```

The argument *local\_rvgname* is the name of the RVG on the local host. The local RVG name represents its RDS.

To display additional information about an RDS, type:

```
# vradmin -g diskgroup -l printrvg local_rvgname
```

The `-l` option with the `printrvg` command displays additional information about the RVGs in an RDS in a long format, such as the data volume count, the SRL name, and the RLINK names for each RVG in the RDS. This `-l` option also displays any configuration errors.



To display information about all the RDSs in the specified disk group, type:

```
# vradmin -g diskgroup printrvg
```

To display information about the RDSs that have local RVGs of name *local\_rvgname*, type:

```
# vradmin printrvg local_rvgname
```

## Displaying Data Volumes in a Replicated Data Set

The `vradmin printvol` command displays information about the data volumes in an RDS and can be entered from any host in the RDS.

To display information about the data volumes in an RDS, type:

```
# vradmin -g diskgroup printvol local_rvgname
```

The argument *local\_rvgname* is the name of the RVG on the local host. The local RVG name represents its RDS.

## Using the vxprint Display Commands

This section describes the VVR commands that you can use to view the status of the objects that take part in replication.

### Displaying a Specific RLINK

Use the `vxprint -Pl` command to display detailed information about the status of an RLINK. This command prints one record per RLINK. The following table lists the information displayed in the output.

To display detailed information about a specific RLINK:

```
# vxprint -g diskgroup -Pl rlink_name
```

Disk Group	Name of the disk group.
RLINK Name	Name of the RLINK.
Info	timeout, packet_size, record_id, latency high, low marks and bandwidth_limit.
State	Displays utility state of the RLINK - ACTIVE, STALE, etc.
synchronous, latencyprot, and srlprot	The current configuration settings for the replication mode, the latency mode, and SRL protection.
assoc	The RVG to which the RLINK is associated.
remote_host/IP_addr/port	The name of the remote host for the RLINK, its IP address, and the port number.
remote_dg	The name of the disk group on the remote system.
remote_dg_dgid	The disk group ID assigned when the disk group was created on the remote system.
remote_rvg_version	The rvg_version of the remote RVG.
remote_rlink	The name of the corresponding RLINK on the remote host.
remote_rlink_rid	The record_id of the corresponding RLINK on the remote host.



local_host/IP_addr/port	The name of the local host, its IP address, and the port number it uses for communication to the remote host.
protocol	The transport protocol used by the RLINK for communicating between the hosts. The protocol can be either UDP/IP or TCP/IP.
checkpoint	Displays the checkpoint only if the Primary RLINK is attached using the checkpoint or the Secondary RLINK has been restored using the checkpoint.
flags	See “ <a href="#">Interpreting RLINK Flag Settings</a> ” on page 86.

---

**Note** To display all the RLINK names in a single column, issue the following command:  
`# vxrvg -g diskgroup [-1] getrlinks rvg_name`

---

## Interpreting RLINK Flag Settings

The following table lists the various flags that can appear in the flags field of the `vxprint -P1` output.

---

**Note** The Primary and Secondary RLINKs are communicating only when the `connected` flag is on. However, replication is taking place only if the following set of flags is displayed:

`write enabled attached consistent connected`

In all other cases, a corrective action may be needed.

---

autosync	The RDS is in the process of automatic synchronization.
attached	The RLINK is attached to the RVG.
can_sync	If the <code>inconsistent</code> and <code>can_sync</code> flags are set, there is enough information in the Secondary SRL volume to make the Secondary consistent again and capable of taking over.
cant_sync	The RLINK is inconsistent, and this Secondary needs a complete resynchronization before it can take over or replicate.
connected	The RLINK is connected to the corresponding RLINK on the remote host and replication can take place.

consistent	The state of the data volumes on the Secondary is suitable for takeover.
dcm_logging	DCM is in use, due to either autosync, failback sync, or an SRL overflow.
detached	The RLINK is STALE and not taking part in replication.
disabled	The RLINK is not attached and is not replicating.
disconnected	The two RLINKs are not connected and are not replicating.
enabled	The RLINK is attached. If the <code>connected</code> flag is displayed, replication can take place. If the <code>disconnected</code> flag is displayed, replication is not taking place.
fail	An I/O error occurred while writing to a data volume on the Secondary.
inconsistent	The data in the Secondary volumes is not consistent and the Secondary cannot take over.
needs_recovery	State after an import or reboot. The <code>vxrecover</code> command clears this state.
primary_paused	The Primary RLINK has been paused and the RLINKs are not replicating.
resync_started	The resynchronization of the Secondary has been started.
resync_paused	The resynchronization has been started but is not currently active because of some problem.
secondary_config_err	There is a mismatch between the configuration of the volumes on the Primary and the Secondary - either a volume is missing on the Secondary or its length is not the same as that of the corresponding volume on the Primary.
secondary_log_err	An I/O error has occurred on the Secondary SRL; replication cannot continue until the SRL has been dissociated and a new one associated.
secondary_paused	The Secondary RLINK has been paused and the RLINKs are not replicating.



## Displaying an Individual RVG

The `vxprint -v1` command displays detailed information about the status of an individual RVG. This command is useful to determine the role of the Primary or Secondary RVG and the state of the RVG as seen by the operating system.

To display detailed information about an RVG:

```
# vxprint -g diskgroup -v1 rvg_name
```

The following table lists the output of the `vxprint -v1` command.

Disk Group	Name of the disk group in which this RVG resides.
RVG	Name of the RVG.
info	Displays the <code>last_tag</code> , <code>record_id</code> , configuration version, and the version of the RVG.
state	The current utility and kernel states.
assoc	Data volumes, SRL, and RLINKs associated with the RVG.
att	The RLINKs that are attached. A Primary can have multiple associated, attached RLINKs. A Secondary can have multiple associated RLINKs, but only one attached RLINK.
flags	See <a href="#">“Interpreting RVG Flag Settings”</a> on page 89





## Interpreting RVG Flag Settings

open	This flag is always set on the Secondary when it is waiting to receive an IBC from the Primary.
closed	This flag is always set.
primary/secondary	Indicates the role of the RVG.
enabled/attached	I/O and IOCTLs can be performed.
disabled/detached	I/O and IOCTLs cannot be performed.

## Displaying an Individual Data Volume

Use the `vxprint -l volume_name` command to display information about a specific data volume. The output fields of special interest for VVR are shown in the next table.

assoc	Shows the RVG to which this data volume is associated (rvg=).
logging	Shows the logging type, which should always be DCM for a data volume.

---

**Note** To list all data volumes in an RVG in a single column, issue the following command:  
**# vxrvlg -g diskgroup [-1] getdatavols rvg\_name**

---

## Displaying the Status of a Secondary

Use the `vxrlink status` command to determine the status of a Secondary. This command displays different information depending on what state the replication is in for that Secondary. For example, whether the Primary is currently replicating to the Secondary, synchronizing the Secondary with a checkpoint, using the DCM to resynchronize the Secondary, or using automatic synchronization for the Secondary. To determine the state of the replication, see [“Displaying Consolidated Replication Status”](#) on page 103.

If the state is replicating, the `vxrlink status` command displays whether the Secondary corresponding to the specified RLINK is up-to-date and if not, how much the Secondary is behind.



Note that outstanding writes are shown even if the Secondary is replicating in synchronous mode. Although for synchronous mode, the write is considered complete to the application when the network acknowledgement is received from the Secondary, VVR still considers the write outstanding until it is written to the data volume on the Secondary.

If automatic synchronization or DCM resynchronization is in progress, the `vxrlink` status command shows the progress of the automatic synchronization.

### ▼ To show the status of a Secondary

```
# vxrlink -g diskgroup status rlink_name
```

If replication is in progress, the output resembles:

```
VxVM VVR vxrlink INFO V-5-1-4640 Rlink rlink_name has <x>
outstanding writes, occupying <y> Kbytes (17%) on the SRL
```

If automatic synchronization is in progress, the output resembles:

```
VxVM VVR vxrlink INFO V-5-1-4464 Rlink rlink_name is in AUTOSYNC.
100864 Kbytes remaining.
```

If DCM resynchronization is in progress, the output resembles:

```
VxVM VVR vxrlink INFO V-5-1-4348 DCM is in use on Rlink rlink_name.
DCM contains 88832 Kbytes.
```

To display the Secondary status periodically, specify a time interval using the `-i` option. For example, to print the status of the Secondary every five seconds, use the command:

```
# vxrlink -g diskgroup -i5 status rlink_name
```

If replication is in progress, the output resembles:

```
VxVM VVR vxrlink INFO V-5-1-4640 Rlink rlink_name has <x>
outstanding writes, occupying <y> Kbytes (17%) on the SRL
VxVM VVR vxrlink INFO V-5-1-4640 Rlink rlink_name has <x>
outstanding writes, occupying <y> Kbytes (19%) on the SRL
```

If automatic synchronization is in progress, the output resembles:

```
VxVM VVR vxrlink INFO V-5-1-4464 Rlink rlink_name is in AUTOSYNC.
100864 Kbytes remaining.
VxVM VVR vxrlink INFO V-5-1-4464 Rlink rlink_name is in AUTOSYNC.
94464 Kbytes remaining.
```

To display the status of an RLINK with a timestamp, use the `vxrlink` status command with the `-T` option. This is useful if the Secondary is not up-to-date. The output of the command displays a timestamp in the locale's appropriate time format to indicate the time by which the Secondary is behind.



For example, if there are pending writes in the Primary SRL, use the following command to check the status of the Primary:

```
# vxrlink -g diskgroup -T status rlink_name
```

The output resembles:

```
VxVM VVR vxrlink INFO V-5-1-4640 Rlink rlink_name has <x> outstanding
writes, occupying <y> Kbytes (20%) on the SRL
VxVM VVR vxrlink INFO V-5-1-0 Rlink rlink_name is behind by 0:00:40
hours
```

The second message indicates the time by which the RLINK is behind.

---

**Note** If the system time is reset to a value different from that of the current system time, then, the output of the `vxrlink -T status` command will appropriately show a negative or an inaccurate value, until the updates that were done before resetting the system time get replicated.

---

## Displaying a List of Checkpoints

VVR enables you to get a list of the Primary checkpoints using the `vxrvrg cplist` command; the `vxrlink cplist` command enables you to get a list of the Secondary checkpoints. The `vxrvrg cplist` and `vxrlink cplist` commands can be run on the Primary only. VVR supports a maximum of 46 checkpoints, and hence the list displays a maximum of 46 checkpoints. If you try to create more than the specified number of checkpoints, an error message prompts you to delete older checkpoints before creating new checkpoints.

Primary checkpoints are created using the `vxrvrg -c checkpoint_name checkstart` command on the Primary and are associated with an RVG. Issue the `vxrvrg cplist` command to display a list of the existing Primary checkpoints associated with the specified RVG. The Primary checkpoints can be deleted using the `vxrvrg -c checkpoint_name checkdelete rvg_name` command.

Secondary checkpoints are created using the `vxrlink -c checkpoint_name pause` command on the Secondary and are associated with the RLINK. Issue the `vxrlink cplist` command on the Primary to display a list of the existing Secondary checkpoints associated with the specified RLINK. The Secondary checkpoint can be deleted by using the `vxrlink -c checkpoint_name checkdelete rlink_name` command.

---

**Note** The `vxrlink cplist` command and the `vxrlink checkdelete` command must be run on the Primary only.

---



The displayed information includes details about each checkpoint, such as checkpoint name, size, percentage of SRL used, and whether the checkpoint has been started or completed. If the SRL usage reaches 100%, the checkpoint overflows and becomes unusable. In this case, VVR displays the message `Checkpoint overflowed`.

To display a list of Primary checkpoints, enter the following command on the Primary:

```
# vxrvg -g diskgroup cplist rvg_name
```

To display a list of Secondary checkpoints, enter the following command on the Primary:

```
# vxrlink -g diskgroup cplist rlink_name
```

where *rlink\_name* is the name of the Primary RLINK that connects to the Secondary where the `vxrlink -c checkpoint_name pause` was issued. The output resembles:

Name	MBytes	% Log	Started/Completed
----	-----	-----	-----
a8	200	5	Completed
a9	800	20	Completed
a6	2000	40	Started

## Using the vrstat Display Commands

This section describes the VVR commands that you can use to view statistics about the RLINKs and the volumes in an RVG, for all the hosts in an RDS. The `vrstat` command combines the output of commands such as `vxrlink stats`, `vxrlink status`, `vxstat`, and `vxmemstat` in a single command to display the statistics about the RLINKs and the volumes in an RVG, for all hosts in the RDS.

The messages are displayed at a default frequency of 10 seconds, which is the frequency at which the `vrstat` command collects the statistics. To change the frequency of the display, set the `VRAS_STATS_FREQUENCY` environment variable to a required value in the `/etc/vx/vras/vras_env` file.

After setting the environment variable to a new value, restart the `vradmin` daemon by issuing the following commands:

```
# /sbin/init.d/vras-vradmind.sh stop  
# /sbin/init.d/vras-vradmind.sh start
```

## Displaying the Consolidated Statistics

To display the consolidated statistics of the RLINKs, SRL, data volumes, and memory tunables for the RDSs on a host, use the `vrstat` command without specifying any option.

To view the consolidated statistics:

```
# vrstat
```



## Displaying the RLINK Information for All the Hosts in the RDS

The `vrstat -R` command displays detailed statistics for the RLINKs on all the hosts in an RDS. This information can be used to assess connectivity and network problems between the hosts. The `vrstat -R` command can be executed from the Primary and the Secondary. The output of this command is a combination of the outputs of the `vxlink stats` and `vxlink status` commands.

To view information about all the RLINKs in an RDS:

```
# vrstat -R [local_rvgname]
```

The argument `local_rvgname` is the name of the RVG on the local host and is optional. The local RVG name represents its RDS.

If you specify the `local_rvgname` name, the `vrstat -R` command displays the information for all the RLINKs for the specified RVG. Otherwise, the command displays the information for all the RLINKs for all the RDSs.

The output of this command resembles:

```
Mon Oct 27 15:44:21 2003
Replicated Data Set hr_rvg:

Data Status:
london: up-to-date.

Network Statistics:
      Messages              Errors              Flow Control
      -----              -
#  Blocks  RT(msec)  Timeout  Stream  Memory Delays  NW Bytes  NW Delay  Timeout
seattle - london
260      133120   5         1         0         0         333      178000   1         20
279         0       11         0         0         0         0       100000   1         30
```

For more information on the fields in the output, see [“Getting and Using Network Performance Data \(vxlink stats\)”](#) on page 111.



## Displaying Information About All the Data Volumes for All the Hosts in the RDS

The `vrstat -V` command displays detailed statistics for all the data volumes associated with the specified RVG on each host in the RDS. The `vrstat -V` command can be executed from the Primary and the Secondary.

To view information about all the data volumes associated with an RVG in an RDS:

```
# vrstat -V [local_rvgname]
```

The argument *local\_rvgname* is the name of the RVG on the local host and is optional. The local RVG name represents its RDS.

If you specify the *local\_rvgname* name, the `vrstat -V` command displays information about all the volumes associated with the specified RVG. Otherwise, the command displays information about all the volumes in all the RDSs.

The output of this command resembles:

```
Mon Oct 27 15:49:15 2003
Replicated Data Set hr_rvg:
```

```
Data Volume-I/O Statistics:
```

HOST	NAME	OPERATIONS		BLOCKS		AVG TIME(ms)	
		READ	WRITE	READ	WRITE	READ	WRITE
seattle	hr_dv01	0	0	0	0	0.0	0.0
london	hr_dv01	0	412	0	210944	0.0	12.0

```
Mon Oct 27 15:49:25 2003
Replicated Data Set hr_rvg:
```

```
Data Volume-I/O Statistics:
```

HOST	NAME	OPERATIONS		BLOCKS		AVG TIME(ms)	
		READ	WRITE	READ	WRITE	READ	WRITE
seattle	hr_dv01	0	0	0	0	0.0	0.0
london	hr_dv01	0	0	0	0	0.0	0.0

The output of this command includes the following details:

- ◆ The *host\_name* of the host for which the information is being displayed.
- ◆ Name of the volume for which the information is being displayed.
- ◆ The total number of read and write operations performed on a volume.
- ◆ Number of blocks that have been read from or written to the volume.
- ◆ Average time in milliseconds to complete the read and write operations.



## Displaying Information About the SRL Volumes for All the Hosts in the RDS

The `vrstat -S` command displays detailed statistics about the SRL for every host within the RDS. This command can be executed from the Primary and the Secondary.

To view information on all the SRL volumes for every host in an RDS:

```
# vrstat -S [local_rvgname]
```

The argument *local\_rvgname* is the name of the RVG on the local host and is optional. The local RVG name represents its RDS.

If you specify the *local\_rvgname* name, the `vrstat -S` command displays information about the SRL in the RDS. Otherwise, the command displays information about the SRL in all the RDSs.

The output of the `vrstat` command resembles:

```
Mon Oct 27 15:53:11 2003
Replicated Data Set hr_rvg:
```

SRL-I/O Statistics:

HOST	NAME	OPERATIONS		BLOCKS		AVG TIME (ms)	
		READ	WRITE	READ	WRITE	READ	WRITE
seattle	hr_srl	0	258	0	1328850.0	17.6	
london	hr_srl	0	0	0	00.0	0.0	

```
Mon Oct 27 15:53:21 2003
Replicated Data Set hr_rvg:
```

SRL-I/O Statistics:

HOST	NAME	OPERATIONS		BLOCKS		AVG TIME (ms)	
		READ	WRITE	READ	WRITE	READ	WRITE
seattle	hr_srl	0	143	0	73430	0.0	17.6
london	hr_srl	0	0	0	0	0.0	0.0

The output of this command includes the following details:

- ◆ The *host\_name* of the host for which the information is being displayed.
- ◆ Name of the SRL volume for which the information is being displayed.





- ◆ The total number of read and write operations performed on a volume.
- ◆ Number of blocks that have been read from or written to the volume.
- ◆ Average time in milliseconds to complete the read and write operations.

## Displaying Information About the Memory Tunable Parameters for All the Hosts in the RDS

The `vrstat -M` command displays detailed information about the memory tunable parameters. This command can be executed from the Primary and the Secondary. The output of the `vrstat -M` command is similar to the output displayed by the `vxmemstat` command.

If you specify the `local_rvgname` name with the `vrstat -M` command, it displays the information about the memory tunables for all the hosts in that RDS. Otherwise, the command displays information about the memory tunable parameters for all the hosts in all the RDSs.

To view information about the memory tunable parameters:

```
# vrstat -M [local_rvgname]
```

The argument `local_rvgname` is the name of the RVG on the local host and is optional. The local RVG name represents its RDS.

The output of this command resembles:

```
Mon Oct 27 15:57:15 2003
Replicated Data Set hr_rvg:
```

```
Memory-pool Statistics:
```

Host	Pool	DG	Min Size	Max Size	In Use	Allocated	Max Used	Waiting
seattle	Voliomem	-	1024	12443	0	1024	0	no
seattle	NMCOM-hr_rvg	dg1	1024	4096	0	1024	0	no
seattle	RViomem	-	1024	12443	0	1024	0	no
seattle	WRSHIP	-	1024	4096	0	1024	0	no
seattle	Readback	-	1024	4096	0	1024	0	no
london	Voliomem	-	1024	12441	0	1024	0	no
london	NMCOM-hr_rvg	dg1	1024	4096	0	1024	0	no
london	RViomem	-	1024	12441	0	1024	0	no
london	WRSHIP	-	1024	4096	0	1024	0	no
london	Readback	-	1024	4096	0	1024	0	no

The output of the `vrstat` command includes the following details:

- ◆ The `host_name` of the host for which the information is being displayed.
- ◆ Name of the memory tunable parameter.



- ◆ Name of the disk group in which this RVG is present.
- ◆ Minimum and maximum size for each tunable parameter.
- ◆ The amount of the allocated space that is being used.
- ◆ Amount of space allocated for the parameter.
- ◆ Maximum space that has been used by a parameter.

## Determining the Network Bandwidth Being Used by VVR

Use the `vrstat` command to determine the network bandwidth being used by VVR.

### ▼ To view the network bandwidth currently being used by VVR

```
# vrstat -R local_rvgname
```

The argument `local_rvgname` is the name of the RVG on the local host and is optional. The local RVG name represents its RDS.

If you specify the `local_rvgname` name, the `vrstat -R` command displays the information about the RLINKs for the specified RVG. Otherwise, the command displays the information for all the RLINKs for all the RDSs.

### Example:

To view the network bandwidth used by the RDS `hr_rvg` between the Primary `seattle` and the Secondary `london`, issue the following command on any host in the RDS:

```
# vrstat -R hr_rvg
```

The output resembles:

```
Replicated Data Set hr_rvg:
```

```
Data Status:
```

```
london: DCM contains 1157888 Kbytes.
```

```
Network Statistics:
```

Messages			Errors			Flow Control		
Blocks	RT(msec)	Timeout	Stream	Memory	Delays	NW Bytes	NW Delay	Timeout
seattle - london								
356	182272	6	1	0	0	280	271000	1 10
339	0	15	0	0	0	0	100000	1 20

Bandwidth Utilization 72908 Kbps.

## Collecting Statistics of the VVR Components

You can configure VVR to collect statistics of the VVR components. The collected statistics can be used to monitor the system and diagnose problems with the VVR setup. VVR collects the statistics generated by the VVR commands `vxlink stats`, `vxlink status` and `vxrvg stats` for all the imported disk groups, and the system level commands `netstat`, `vmstat`, and `vxmemstat`. The output of these commands are stored in separate files.

By default, VVR collects the statistics automatically when the `vradmind` daemon starts. Configuring VVR to collect statistics according to your requirements involves modifying the values of the environment variables in the `vras_env` file, located in the `/etc/vx/vras` directory.

---

**Note** If the `vradmind` daemon is not running, VVR stops collecting the statistics.

---



**▼ To configure VVR to collect statistics automatically**

1. Modify the default values for the environment variables specified in the `vras_env` located in the `/etc/vx/vras` directory file to suit your requirements. The following table provides information about the variables:

---

<b>Environment Variable</b>	<b>Description</b>
<code>VRAS_ENABLE_STATS</code>	Specifies whether you want the statistics collection to start automatically. Set <code>VRAS_ENABLE_STATS=on</code> to enable statistics collection. This is the default. Set <code>VRAS_ENABLE_STATS=off</code> to disable statistics collection.
<code>VRAS_STATS_FREQUENCY</code>	Specifies the frequency in seconds at which the statistics should be collected for the VVR commands, <code>vxrlink stats</code> , <code>vxrlink status</code> and <code>vxrvrg stats</code> . By default, <code>VRAS_STATS_FREQUENCY</code> is set to 10 seconds.
<code>VRAS_NETSTAT_FREQUENCY</code>	Specifies the time interval in seconds over which the statistics for the different network protocols should be collected. By default, <code>VRAS_NETSTAT_FREQUENCY</code> is set to 300 seconds.
<code>VRAS_VMSTAT_FREQUENCY</code>	Specifies the time interval in seconds over which the memory and CPU utilization statistics should be collected. By default, <code>VRAS_VMSTAT_FREQUENCY</code> is set to 300 seconds.
<code>VRAS_STATS_DAYS_LOG</code>	Specifies the number of days for which the collected statistics should be preserved. After this time the earlier statistics are automatically deleted. By default, <code>VRAS_STATS_DAYS_LOG</code> is set to three days.

---

2. Restart the `vradmin` daemon by issuing the following command:

```
# /sbin/init.d/vras-vradmin.sh stop  
# /sbin/init.d/vras-vradmin.sh start
```

You can restart `vradmin` even while the application is active and replication is in progress.

## Understanding How VVR Stores the Statistics

VVR stores the statistics collected for each of the commands, that is, `vxmemstat`, `netstat`, `vmstat`, `vxrlink stats`, `vxrlink status` and `vxrvlg stats` in separate files that are stored in the `/var/vx/vras/stats/` directory. Each file stores the statistics for a day and only those files for the period specified by the `VRAS_STATS_DAYS_LOG` variable are preserved. The earlier files are automatically deleted. VVR stores the statistics in files that are named using the following convention:

- ◆ `statsType_dgName_objectN`—stores the statistics of the VVR components collected by the commands `vxrlink stats`, `vxrlink status` and `vxrvlg stats`.
- ◆ `statsType_hostname_date`—stores the system level statistics collected by the commands `netstat`, `vmstat`, and `vxmemstat`.

The data collected for the `vxmemstat` command is the same as that displayed by the `vxmemstat -e` command.

The output collected for the `vxrlink stats` command is a combination of the fields displayed by the `vxrlink -e stats` command and the `vxrlink stats` command. Using network performance data the output of `vxrlink stats` is displayed under the following headings. For more information, see [“Getting and Using Network Performance Data \(vxrlink stats\)”](#) on page 111.

### Messages

The fields displayed under this heading are the same as those displayed by the `vxrlink stats` command, with the exception of an additional field `Blocks (ACKed)`. This field displays information about the number of transmitted blocks that have been acknowledged.



## Errors

Most of the fields under this heading are similar to those displayed by the `vxrlink -e stats` command; however, some of the field names are different. The following table provides the mapping:

---

<b>Fields in the vxrlink stats command when vradmind collects the statistics</b>	<b>Fields in the vxrlink -e stats command</b>
Memory	No memory available
Slots	No message slots available
Pool	No memory available in nmcom pool on Secondary
Timeout	Timeout
Packet	Missing packet
Message	Missing message
Stream	Stream
Checksum	Checksum
Transaction	Unable to deliver due to transaction

---

## Flow control

The fields displayed under this heading are the same as those displayed by the `vxrlink stats` command. For more information on the `vxrlink stats` command, see [“Getting and Using Network Performance Data \(vxrlink stats\)”](#) on page 111.



## Displaying Consolidated Replication Status

The `vradmin repstatus` command displays the consolidated replication status of the specified Replicated Data Set (RDS). The `vradmin repstatus` command displays the following information about each RVG in the RDS:

- ◆ Consolidated view of the RDS
- ◆ Replication settings for all Secondary hosts in the RDS
- ◆ Status of the data on each Secondary host in the RDS
- ◆ Status of replication to each Secondary host in the RDS

To display consolidated replication information about an RDS:

```
# vradmin -g diskgroup [-l] repstatus local_rvgname
```

The argument `local_rvgname` is the name of the RVG on the local host. The local RVG name represents its RDS.

The option `-l` displays additional information, such as RLINK names, replication setting, and so on. Similar to the `vradmin -l printrvg` command, the `vradmin repstatus` command also displays configuration errors in the RDS, if any.

---

**Note** If the `vradmin repstatus` command is run on a Secondary that has a few configuration errors or cannot reach the Primary, the output displays the status known to the Secondary before the above condition occurred and therefore might be out of date.

---



**Example:**

When the Primary is reachable from all the Secondary hosts and the `vradmin repstatus` command is run from any host:

```
# vradmin -g hrdg -l repstatus hr_rvg
```

Output resembles:

Replicated Data Set: hr\_rvg

Primary:

```
Host name:          seattle
RVG name:           hr_rvg
DG name:            hrdg
RVG state:          enabled for I/O
Data volumes:       2
SRL name:           hr_srl
SRL size:           4.00 GB
Total secondaries: 1
```

Secondary:

```
Host name:          london
RVG name:           hr_rvg
DG name:            hrdg
Rlink from Primary: rlk_london_hr_rvg
Rlink to Primary:   rlk_seattle_hr_rvg
Configured mode:    asynchronous
Latency protection: off
SRL protection:     autodcm
Data status:         inconsistent
Replication status: resync in progress (autosync)
Current mode:        asynchronous
Logging to:          DCM (contains 169728 Kbytes) (autosync)
Timestamp Information: N/A
Bandwidth Limit:    30.00 Mbps
```





**Example:**

When the Primary is unreachable from the Secondary hosts and the `vradmin repstatus` command is run from the Secondary host:

```
# vradmin -g hrdg -l repstatus hr_rvg
```

Output resembles:

VxVM VVR vradmin INFO V-5-52-1205 Primary is unreachable or RDS has configuration error. Displayed status information is from Secondary and can be out-of-date.

Replicated Data Set: hr\_rvg

Primary:

```
Host name:          seattle <unreachable>
RVG name:           hr_rvg
DG name:            hrdg
RVG state:          enabled for I/O
Data volumes:       2
SRL name:           hr_srl
SRL size:           4.00 GB
Total secondaries: 1
```

Secondary:

```
Host name:          london
RVG name:           hr_rvg
DG name:            hrdg
Rlink from Primary: rlk_london_hr_rvg
Rlink to Primary:   rlk_seattle_hr_rvg
Configured mode:    asynchronous
Latency protection: off
SRL protection:     autodcm
Data status:         consistent, up-to-date
Replication status: replicating (connected)
Current mode:        asynchronous
Logging to:          SRL (0 updates behind, last update ID 18533.0)
Timestamp Information: behind by 00:00:00 hours
Bandwidth Limit:    30.00 Mbps
Last Update on Primary: Oct 10 04:32:21
Secondary up-to-date as of: Oct 10 04:32:21
```

Config Errors:

```
seattle:           Pri or Sec IP not available or vradmind not running,
                   stale information
```



The following section describes the important fields displayed by the `vradmin repstatus` command. The values and meaning of each field are listed in tables:

- ◆ **RVG state:** Displays the state of the Primary RVG. The following table lists the values for the RVG state field and their meanings.

Value	Meaning
acting_secondary	This Primary RVG is currently the acting Secondary as part of the fast failback process. Writes to the data volumes in this RVG are disabled independent of whether the RVG is started or stopped.
disabled for I/O	Primary RVG is disabled for I/O, that is, the RVG is stopped.
enabled for I/O	Primary RVG is enabled for I/O, that is, RVG is started.
needs recovery	State after an import or reboot. The <code>vxrvrg recover rvg</code> command clears this state.
passthru	The Primary RVG is in <code>passthru</code> mode because the Primary SRL is detached or missing. For more information, see “ <a href="#">RVG PASSTHRU Mode</a> ” on page 315.

- ◆ **Data status:** Shows the data status of this Secondary. The following table lists the values for the Data status field and their meanings:

Value	Meaning
consistent, behind	Secondary data is consistent but not up-to-date with the Primary data.
consistent, stale	The data on this Secondary is consistent. Replication to this Secondary has been stopped; the Primary RLINK is detached.
consistent, up-to-date	The Secondary data is consistent and is current or up-to-date with the Primary data. The Primary role can be migrated to this Secondary.

inconsistent	The data on the Secondary volumes is not consistent and the Secondary cannot take over.
needs recovery	State after an import or reboot. The <code>vxrlink recover</code> command clears this state.
N/A	Current state of the Secondary data cannot be determined. This may occur because of a configuration error on this Secondary. For information about the state, use the <code>vxprint -l rlink_name</code> command on the Primary and Secondary.

- ◆ **Current mode:** Displays the mode of replication, asynchronous or synchronous, that is being used to replicate data to the Secondary. This value can be different from the configured replication setting if the configured mode is `synchronous=override`. For more information, see [“Mode of Replication—synchronous attribute”](#) on page 27.
- ◆ **Replication status:** Displays the status of the replication to the Secondary. The following table lists the values for the Replication status field and their meanings:

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 <a href="#">“Incrementally Synchronizing the Secondary After SRL Overflow”</a> on page 135.  needs failback synchronization—To continue replication, start failback synchronization to this Secondary. See <a href="#">“Failing Back Using Fast Failback Synchronization”</a> on page 232.
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 <code>needs_recovery</code> state.  primary <code>needs_recovery</code> —Primary RLINK in <code>needs_recovery</code> state and needs to be recovered before replication can resume.
paused by user	Replication to Secondary is paused because of some administrative action. This results in the following states:  primary paused—Primary RLINK is paused. secondary paused—Secondary RLINK is paused.
paused due to error	Replication to Secondary is paused because of the following errors:  secondary config error—Secondary has some configuration error. For more information, refer to the <code>secondary_config_err</code> in the section <a href="#">“Interpreting RLINK Flag Settings”</a> on page 86. secondary log error—Secondary SRL has an I/O error. See the <code>secondary_log_err</code> in section <a href="#">“Interpreting RLINK Flag Settings”</a> on page 86.
paused due to network disconnection	Replication to Secondary is paused because of some network problem.
replicating	connected—Replication can take place if there are updates on the Primary data volumes
resync in progress	Resynchronization to the Secondary is in progress. autosync—Resynchronization type is autosync. dcm resynchronization—Resynchronization after an SRL overflow. failback resynchronization—Resynchronization using failback logging.



resync paused by user	<p>Resynchronization to Secondary is paused because of some administrative action. This results in the following states:</p> <p>primary paused—Primary RLINK is paused.</p> <p>secondary paused—Secondary RLINK is paused.</p>
resync paused due to error	<p>Resynchronization to Secondary is paused because of the following errors:</p> <p>secondary config error—Secondary has some configuration error. For more information, refer to the <code>secondary_config_err</code> in the section <a href="#">“Interpreting RLINK Flag Settings”</a> on page 86.</p> <p>secondary log error—Secondary SRL has an I/O error. See the <code>secondary_log_err</code> in section <a href="#">“Interpreting RLINK Flag Settings”</a> on page 86.</p>
resync paused due to network disconnection	<p>Resynchronization to Secondary is paused because of some network problem.</p>
stopped	<p>Replication to Secondary is stopped because of the following:</p> <p>Primary detached—Primary RLINK is detached.</p> <p>Secondary detached—Secondary RLINK is detached.</p>
N/A	<p>The replication status cannot be determined. For information about the status, use the <code>vxprint -l rlink_name</code> command on the Primary and Secondary.</p>



- ◆ **Logging to:** Indicates whether updates for this Secondary are tracked on the Primary using the SRL or DCM. The following table lists the values for the Logging to field and their meanings:

Value	Meaning
DCM (contains xxx Kbytes) (log_type)	DCM is active (in use) for the replication to this Secondary. log_type can be autosync, failback logging, or SRL protection logging.
SRL (xxx Kbytes behind, yyy % full)	Updates to be transferred to Secondary are logged into the SRL and are currently occupying xxx Kbytes or yyy% of the SRL
SRL	SRL is used for logging. Check the Data status field for the status of the Secondary data.

If the `vradmin repstatus` command is run on a Secondary and the Secondary is disconnected from the Primary because of a configuration or network error, the Logging to field may show the following values:

Value	Meaning
DCM (log_type)	The last known information about logging type before the Secondary disconnected from the Primary is that it was logging to DCM.
SRL (xxx updates behind, last update ID yyy)	Before Secondary disconnected from the Primary, SRL was used for logging. Secondary was xxx updates behind the Primary, and the last update that was applied on the Secondary has update ID yyy. This information is similar to that displayed by the <code>vxrlink updates</code> command.
SRL (updates behind N/A)	Before Secondary disconnected from the Primary, SRL was used for logging. The number of updates this Secondary is behind with the Primary is not known.

- ◆ **Timestamp information, Last Update on Primary, Secondary up-to-date as of:**

These fields are the same as the output that is displayed for the `vxrlink -T` command. For more information, see [“Displaying the Status of a Secondary”](#) on page 89 and [“Identifying the Most Up-to-Date Secondary”](#) on page 114.



## Getting and Using Network Performance Data (vxlink stats)

The `vxlink stats` command reports detailed information about the state of the network. It displays network statistics that can be used to assess network problems. Use the network performance data to determine the optimum network configuration for efficient use of system resources.

The `vxlink stats` command can be executed repeatedly at given intervals using the `-i` interval option. In this case, the displayed values indicate the change since the last interval, except for the average round-trip value, which displays a moving average. The `vxlink stats` command can be executed from the Primary and the Secondary. The RLINK statistics are reset when the RLINK disconnects. For detailed information on available options, refer to the `vxlink` manual page.

The output of the `vxlink stats` command includes the following details:

- ◆ Number of messages transmitted.
- ◆ Number of 1024-byte blocks transmitted
- ◆ Average round-trip per message.  
The size of the message affects the average round-trip per message.
- ◆ Number of timeouts or lost packets.  
If the number of timeouts is high, it indicates that network is very lossy. This needs to be fixed.
- ◆ Number of stream errors  
Stream errors occur when the RLINK attempts to send messages faster than the network can handle.
- ◆ Number of memory errors  
Memory errors occur when the secondary has insufficient buffer space to handle incoming messages. To reduce the number of errors, try increasing the value of the tunable, `vol_max_nmpool_sz` on the secondary.
- ◆ Current Timeout Value  
This represents the packet timeout in milliseconds.



## Displaying Extended Replication Statistics

You can use the `vxlink stats` command with the `-e` option to generate extended statistics, in addition to the statistics generated by the `vxlink stats` command. The output generated by this command can be useful in assessing the reason for failure at the time it occurred.

The `vxlink stats -e` command can be executed repeatedly at given intervals using the `-i` interval option. In this case, the displayed values indicate the change since the last interval. The `vxlink stats -e` command can be executed from the Primary as well as the Secondary. The RLINK statistics are reset when the RLINK disconnects.

For detailed information about the available options, refer to the `vxlink` manual page.

The output of the `vxlink stats -e` command is displayed under the headings Messages and Errors. Each of these headings has the appropriate fields to display the required information. The first is the Messages heading which displays the following information:

- ◆ Number of blocks sent

Displays the number of 1024 bytes that have been transmitted. This is different from the `Blocks` attribute displayed by the `vxlink stats` command (without the `-e` option), which only displays the number of blocks that have been acknowledged.

The Messages heading is followed by the Errors heading. It has nine fields that display the different kinds of error encountered, three of which are similar to that in the `vxlink stats` command. The output includes the following details:

- ◆ No memory available

This error occurs when there is no space in the systems kernel memory to process the message.

- ◆ No message slots available

This error occurs if there is no memory to store the packets of the message that have arrived out of sequence. If a packet arrives out of sequence then it requires to be stored in the message buffer until all the related out-of-sequence packets arrive and can be assembled.

- ◆ No memory available in `nmcom` pool on Secondary

The buffer space determined by the VVR tunable `vol_max_nmpool_sz` is already full and cannot store any new messages that arrive at the Secondary.

- ◆ Timeout errors

Indicates the number of timeout errors, that is, the number of times the Primary timed out while waiting for an acknowledgement from the Secondary.



- ◆ Missing packet errors  
Indicates the number of times the last packet of a message was received, before one or more packets of the same message were received.
- ◆ Missing message errors  
Indicates the number of times messages have arrived out of sequence.
- ◆ Stream errors  
Stream errors occur when the RLINK attempts to send messages faster than the network can handle.
- ◆ Checksum errors  
Displays the number of data checksum errors. Every time a packet is received at the Secondary VVR performs a checksum to ensure that the packet data is the same as that sent by the Primary.
- ◆ Unable to deliver due to transaction errors  
Displays the number of times the packets could not be delivered to the Secondary, due to transaction errors. If the Secondary is busy with some kernel operations when the packet arrives at the Secondary, then these packets may not be delivered until the transaction is complete.



## Identifying the Most Up-to-Date Secondary

VVR provides the `vxlink updates` command to identify the most up-to-date Secondary in a VVR configuration. The `vxlink updates` command can be issued on a Secondary only.

For multiple Secondaries, the `vxlink updates` command enables you to determine the Secondary that contains the most up-to-date data and hence the most suitable replacement for the Primary in the case of a takeover.

For a single Secondary, the `vxlink updates` command can be used to determine the extent to which the Secondary is behind the Primary. You can decide whether or not to take over the Primary role by looking at the update ID of the Secondary, number of updates the Primary is ahead of the Secondary, and how long you expect the Primary to be unavailable.

Issue the following command on the Secondary.

```
# vxlink -g diskgroup -T updates rlink_name
```

To display only the update ID in the output, use the `vxlink updates` command without the `-T` option. The output displays the update ID as a sequence number. A sequence number is a 64-bit value that increases incrementally and therefore is unique for each new update. The output of the `vxlink updates` command displays the 64-bit number as two 32-bit sequence numbers separated by a dot. For example,

```
high_seq_num . low_seq_num
```

To display the exact time on the Primary at which the Secondary is up-to-date, use the `vxlink updates` command with the `-T` option. The `-T` option displays the exact time in hours by which the Secondary is behind. Note that the update information may be inaccurate if:

- ◆ the Secondary has been rebooted and even before it comes up, the Primary becomes unavailable.
- ◆ the Secondary reboots and the RLINK gets disconnected.

The output of the `vxlink -T updates` command is displayed in a three column structure with two rows; ID and Time. The ID row displays the update IDs. The timestamp in the Time row indicates the time at which the update was written on the Primary. The time is displayed in *Mon date time* format, where *Mon* is a locale abbreviated month name followed by the *date* and the *time* in the locale's appropriate time format.

The first column displays the last update ID and the time at which it was written on the Primary.

The second column displays the last update ID that has been received on the Secondary and the time when it was written on the Primary. If the Secondary is up-to-date then the ID and the time in this column is the same as that in the first column. However, if the Secondary is behind, then the ID and the time is different from that in the first column.

The third column indicates the exact number of updates by which the Secondary is behind and also the time in the locale's appropriate time format by which it is behind. This value is obtained as a difference between the second and first column.

---

**Note** If the system time is reset to a value different from that of the current system time, the output of the `vxrlink -T updates` command appropriately shows a negative or an inaccurate value, until the updates that were completed before resetting the system time are replicated.

---

## Example—To Determine the Most Up-to-Date Secondary

This example shows how to determine the most up-to-date Secondary in an RDS containing Primary `seattle` and the Secondaries `london` and `newyork`. This example displays the last update ID received by the Secondary and the last known update ID on the Primary.

1. On the Secondary `london`, enter the following command:

```
# vxrlink -g diskgroup updates to_seattle
```

The output resembles:

```
Secondary has received an update ID of 37364.104, last known update ID on Primary is 99 updates ahead.
```

2. On the Secondary `newyork`, enter the following command:

```
# vxrlink -g diskgroup updates to_seattle
```

The output resembles:

```
Secondary has received an update ID of 37364.118, last known update on Primary is 95 updates ahead.
```

Compare the output on `london` and `newyork`. The host `newyork` has received the later update with the update ID 37364.118, whereas the other Secondary `london` is behind `newyork` by 14 updates. The host `newyork` is more up-to-date than the host `london`.



This example shows how to determine the status of the Secondary in an RDS containing Primary `seattle` and the Secondary `london`, using the `-T` option with the `vxlink updates` command.

1. On the Secondary `london`, enter the following command:

```
# vxlink -g diskgroup -T updates to_seattle
```

If the Secondary is up-to-date the output displays the following:

	Last update	Secondary	Secondary
	on Primary	up-to-date as of	behind by
ID	34666.0	34666.0	0
Time	Oct 16 11:17:44	Oct 16 11:17:44	00:00:00

If the Secondary is not up-to-date the output displays the following:

	Last update	Secondary	Secondary
	on Primary	up-to-date as of	behind by
ID	34666.640	34666.592	48
Time	Oct 16 11:17:44	Oct 16 11:17:42	00:00:02

## VVR Event Notification

VVR provides the `vrnotify` utility to notify administrators of VVR specific events, such as SRL full, resynchronization complete, etc. You can receive notification for a VVR event on the Primary or Secondary node, or both the nodes in an RDS.

The `vrnotify` command enables you to write a script that receives VVR event notification and notifies administrators of these events through email, pager, etc. See the examples in this section to see how event notifications can also be used to keep history of various events.

If you do not specify the `local_rvgname` in the `vrnotify` command, event notification is started for all the RDSs on the local host. Use the `-g` option to receive event notifications for RVGs in a specific disk group.

The `vrnotify` command displays the VVR events until you explicitly terminate or kill the command.

To receive event notifications on the Primary or Secondary, enter the following command:

```
# vrnotify -g diskgroup local_rvgname...
```

The argument `local_rvgname...` is a space-separated list of the names of the RVGs on the local host in the specified disk group.

The `vrnotify` command displays each event on a new line in the following format:

```
host_name: event_type: RDS_name: event message
```

The `vrnotify` command displays the following types of events:

Event Type	Event Message
<code>resync_started</code>	Resync started on Primary RVG
<code>resync_stopped</code>	Resync stopped on Primary RVG
<code>resync_paused</code>	Resync paused on Primary RVG
<code>lat_throttle_on</code>	Latency throttling started
<code>lat_throttle_off</code>	Latency throttling stopped
<code>lat_throttle_override</code>	Latency throttling overridden
<code>lat_throttle_fail</code>	Latency throttling caused I/O failures
<code>srlprot_throttle_on</code>	SRL overflow protection throttling started



<b>Event Type</b>	<b>Event Message</b>
srlprot_throttle_off	SRL overflow protection throttling stopped
srlprot_override	SRL overflow protection overridden
srlprot_fail	SRL overflow protection caused I/O failures
srl_overflow	Replication stopped due to SRL overflow
srlprot_dcm_on	Started using DCM for SRL protection
srlprot_dcm_off	Stopped using DCM
rlk_connect	RLINK connected to remote
rlk_disconnect	RLINK disconnected from remote
srl_log_warn	SRL over 80% full
repmode_sync	Replicating in synchronous mode
repmode_async	Replicating in asynchronous mode
repibc_freeze	Replication on Secondary frozen due to IBC
repibc_unfreeze	Replication on Secondary unfrozen after IBC
rvg_pritosec	RVG role changed from Primary to Secondary
rvg_sectopri	RVG role changed from Secondary to Primary
rvg_pritoactsec	RVG role changed from Primary to acting Secondary
rvg_actsectopri	RVG role changed from acting Secondary to Primary
rlk_paused	Secondary RLINK paused because of a configuration error
ibcmgs_discarded	IBC was discarded due to timeout on the Secondary.



**Example:**

The following example script shows how to use the `vrnotify` utility to receive event notifications for the `hr_rvg` RDS in the `hrdg` disk group and send email to the alias `vvradmin` if the event `srl_warning` occurs.

```
#!/bin/sh
IFS=:
vrnotify -g hrdg hr_rvg | while read host event rvg msg
do
    case $event in
        srl_log_warn)
            (echo "This message is sent by VVR notify mechanism"
             echo "SRL for RVG $rvg on host $host is 80% full"
             ) | mailx -s "VVR SRL Log Warning" vvradmin;;
    esac
done
```







This chapter includes instructions on how to administer data volumes, SRL, replication, and Replicated Data Sets. It also explains how to perform snapshot operations on RVGs.

## Administering Data Volumes

### Associating a Volume to a Replicated Data Set

The `vradmin addvol` command can be used to add a volume to a RDS even when replication is in progress. This command associates a volume to all the RVGs of the RDS. Note that volumes of the same name and same length must exist on all Secondaries and the Primary of the RDS. You must create volumes of the required layout on the Secondary and Primary hosts before issuing the `vradmin addvol` command. If necessary, the `vradmin addvol` command can be used to add a volume to an RDS that only has a Primary RVG. In this case, there are no Secondary volumes.

By default, the `vradmin addvol` command adds DCM logs to the data volumes being added to the RDS, if they have not already been added. If any of the data volumes contains a DRL log, the `vradmin addvol` command removes the DRL log before adding the DCM to the data volume.

The `-nodcm` option with the `vradmin addvol` command adds data volumes to the RDS without adding DCMs to the data volumes. If any of the data volumes has a DRL, the DRL is removed before the data volume is associated with the RVG. If `-nodcm` is issued when any of the RLINKs has `srlprot` set to `dcm` or `autodcm`, and any of the data volumes being added to the RDS does not already have a DCM log, the command will fail.

The `vradmin addvol` command can be run from any host in the RDS. If the `vradmin addvol` command fails on any of the hosts in the RDS during its execution, the volume is not added on any host.



Before adding a volume, the `vradmin addvol` command 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 the section [“Performing Offline Data Verification”](#) on page 187. To synchronize the Primary and Secondary volumes, see [“Synchronizing Volumes on the Local Host and Remote Hosts”](#) on page 124. To skip this confirmation, use the `-s` option with the `vradmin addvol` command. The `-s` option to the `vradmin addvol` command proves useful in scripts.

**Prerequisites for adding a volume to an RDS:**

- ✓ Create volumes of same name and length as the Primary volume on all the hosts in the RDS.
- ✓ Verify that the volumes to be added are inactive.
- ✓ Synchronize the volumes using the `vradmin syncvol` command before adding the volumes to the RDS. Refer to [“Synchronizing Volumes on the Local Host and Remote Hosts”](#) on page 124.

---

**Note** To add a volume to an RDS that has only a Primary RVG, the prerequisites above do not apply.

---

**▼ To add a volume to an RDS:**

```
# vradmin -g diskgroup addvol local_rvgname volume_name
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `volume_name` is the name of the volume to be added to the RDS. Only one volume can be added at a time.

Use the `-nodcm` option when you do not want to add DCMs to the data volumes. By default, DCMs are automatically added.

**Example:**

This example shows how to add an existing volume `hr_dv01` to all RVGs of the RDS. The disk group `hrdg` contains the local RVG `hr_rvg` of the RDS. To add the volume `hr_dv01` to all RVGs in the RDS and automatically add DCMs to the data volumes, type the following command on any host:

```
# vradmin -g hrdg addvol hr_rvg hr_dv01
```

## Verifying the Data on the Primary and Secondary Volumes

The `vradmin syncvol` command when used with the `-verify` option enables you to verify whether the remote volumes and the corresponding local volumes are identical before adding them to an RDS. Use this command when the volumes are not associated with an RVG and the application is inactive (the volumes are not in use). VVR also allows you to verify the data volumes after they have been added to an RDS. For information, refer to “[Verifying the Data on the Secondary](#)” on page 185.

The `vradmin -verify syncvol` command 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. 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.

---

**Note** Remote volumes can be verified with local volumes only if the `/etc/vx/vras/.rdg` file on the remote host contains a local disk group ID entry. Ensure that each disk group ID entry in the `.rdg` file is on a separate line.

---

Note that the order of the volume names in the local and remote volume lists is important. The `vradmin -verify syncvol` command verifies the first volume in the remote volume list with the first volume in the local volume list, and so on. Hence, the number of volumes in the local and remote volume lists must be same. Also, the remote disk group name must be specified if volume names are different on local and remote hosts.

It is recommended that the names of the volumes on the local and remote hosts be the same. However, you can verify volumes with different names on the local and remote hosts using the `vradmin -verify syncvol` command.

### ▼ To verify the difference between the local and remote data volumes

```
# vradmin -g diskgroup -verify syncvol local_vols_list \  
    remote_hostname...
```

The argument `local_vols_list` is a comma-separated list of volumes on the local host. The names of the volumes on the local and remote hosts are assumed to be the same.

The argument `remote_hostname` is a space-separated list of names of the remote hosts on which the volumes to be verified reside. It must be possible for IP to resolve the remote host names.



**Example:**

This example shows how to verify the differences between the remote volumes on host `london` with the local volumes `hr_dv01`, `hr_dv02`, `hr_dv03` in the disk group `hrdg` on the local host `seattle`. The names of the disk group and the volumes on the remote host are the same as names of the disk group and volumes on the local host.

```
# vradmin -g hrdg -verify syncvol hr_dv01,hr_dv02,hr_dv03 london
```

## Synchronizing Volumes on the Local Host and Remote Hosts

The `vradmin syncvol` command 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 `vradmin syncvol` command 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 `vradmin syncvol` command, and then add it to the RDS.

Using the `vradmin syncvol` command, you can synchronize remote volumes with local volumes using one of the following options:

- ◆ Difference-based synchronization
- ◆ Full synchronization

By default, the `vradmin syncvol` command synchronizes the volumes using difference-based 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 `vradmin syncvol` command.

---

**Note** The `vradmin syncvol` command cannot be used for an RVG that has volume-set component volumes associated to it.

---

You can supply a list of volumes to be synchronized. If you choose this method, the order of the volume names in the local and remote volume lists is important. The `vradmin syncvol` command synchronizes the first volume in the remote volume list with the first volume in the local volume list, and so on. Hence, the number of volumes in the local and remote volume lists must be the same. Also, the remote disk group name must be specified if volume names are different on the local and remote hosts.

---

**Note** Remote volumes can be synchronized with local volumes only if the `/etc/vx/vras/.rdg` file on the remote host contains a local disk group ID entry. Ensure that each disk group ID entry in the `.rdg` file is on a separate line.

---

To enable the `vradm` `syncvol` command for a specific disk group on a remote host, enter the local disk group ID in the `/etc/vx/vras/.rdg` file on the remote host. To enable the `vradm` `syncvol` command for all disk groups on a remote host, enter a plus (+) sign in the `/etc/vx/vras/.rdg` file on the remote host. For more information, see the `vradm(1M)` manual page.

Before synchronizing volumes, the `vradm` `syncvol` command displays a warning and prompts the user to confirm whether or not the data on the volumes on the remote host can be overwritten with the data on the volumes on the local host. To skip this confirmation, use the `-s` option with the `vradm` `syncvol` command. The `-s` option to the `vradm` `syncvol` command proves useful in scripts.

### Synchronizing Volumes Using Full Synchronization

In full synchronization, all data is transferred between hosts. Use full synchronization to create initial copies of volumes. To do a full synchronization, specify the option `-full`.

To synchronize volumes on local and remote hosts using full synchronization:

```
# vradm -g diskgroup -full syncvol local_vols_list \  
    remote_hostname...
```

The argument `local_vols_list` is a comma-separated list of volumes on the local host. The names of the volumes on the local and remote hosts are assumed to be the same.

The argument `remote_hostname` is a space-separated list of names of the remote hosts on which the volumes to be resynchronized reside. It must be possible for IP to resolve the remote host names.

#### Example:

This example shows how to do a full synchronization of the remote volumes on host `london` with the local volumes `hr_dv01`, `hr_dv02`, `hr_dv03` in the disk group `hrdg` on the local host `seattle`. The names of the disk group and the volumes on the remote host are the same as the names of the disk group and volumes on the local host.

```
# vradm -g hrdg -full syncvol hr_dv01,hr_dv02,hr_dv03 london
```

### Synchronizing Volumes Using Difference-Based Synchronization

In difference-based synchronization, VVR compares the blocks of data between the hosts and then transfers over the network only those blocks of data that are different. Difference-based synchronization is useful when there is little difference between the data on the local and remote volumes.

To synchronize volumes on local and remote hosts using difference-based synchronization:

```
# vradm -g diskgroup syncvol local_vols_list remote_hostname...
```



The argument *local\_vols\_list* is a comma-separated list of volumes on the local host. In this case, the names of the volumes on the local and remote hosts are the same.

The argument *remote\_hostname* is a space-separated list of names of the remote hosts on which the volumes to be resynchronized reside. It must be possible for IP to resolve the remote host names.

### Example 1:

This example shows how to do a difference-based synchronization of the remote volumes on host *london* with the volumes *hr\_dv01*, *hr\_dv02*, *hr\_dv03* in the disk group *hrdg* on local host *seattle*. The names of the disk group and the volumes on the remote host are the same as names of the disk group and volumes on the local host.

```
# vradmin -g hrdg syncvol hr_dv01,hr_dv02,hr_dv03 london
```

### Example 2:

In this example, the names of the volumes on the remote host are different from the names of the volumes on the local host. It shows how to do a difference-based synchronization of the remote volumes *hr\_dvmaster* and *hr\_dvoralog* on host *london* with the local volumes *hr\_dv01* and *hr\_dv02* in the disk group *hrdg*.

```
# vradmin -g hrdg syncvol hr_dv01,hr_dv02 \  
london:hrdg:hr_dvmaster,hr_dvoralog
```

## Associating the Component Volumes of a Volume Set to an RDS

This section describes how to associate component volumes of a volume set to an RDS. A volume set is a container object for a group of volumes that can be a part of a Multi-device File System (MDFS). Associating the component volumes of a volume set to an RDS enables you to replicate an MDFS.

The component volumes of a volume set have assigned indices. Applications use these indices to identify a component volume. To be able to successfully start applications on the Secondary if a disaster occurs, the component volumes of the volume set on the Secondary must have the same indices as those of the corresponding Primary volumes.

This section assumes that the volume set that is to be replicated already exists on the Primary and its Secondaries. If the volume set does not exist, create the volume set. For more information about volume sets, see the *VERITAS Volume Manager Administrator's Guide*.

---

**Note** You cannot associate a component volume of a volume set as an SRL.

---

Before associating component volumes of a volume set to an RDS, ensure that these volumes have identical indices on the Primary and its Secondaries.

---

**Note** The `vradmin syncrvg` command cannot be used for an RDS that has component volumes of a volume set associated to it. Also, you cannot use the `vradmin syncvol` command to synchronize volumes that are components of a volume set.

---

### ▼ To associate component volumes of a volume set to an RDS

1. Verify whether the component volumes of the volume set on the Primary and its Secondaries have identical indices. To view the indices, use the following command:

```
# vxvset -g diskgroup list vset_name
```

2. If the indices of the component volumes on the Primary volume set and Secondary volume set are identical, go to [step 4](#).
3. If the indices of the component volumes on the Primary volume set and Secondary volume set are different, perform the following steps on the Secondary:

- a. Dissociate each volume from the volume set using the following command:

```
# vxvset -g diskgroup rmvol vset_name compvol_name
```

When you remove the last volume, the volume set is also removed.







## Associating a Data Change Map to a Data Volume

The `vradmin createpri`, `vradmin addsec`, and `vradmin addvol` commands associate the Data Change Map (DCM) to the data volume by default. This section describes how to associate DCMs to a data volume in an existing VVR configuration. For more information about the DCM object, see “[Data Change Map](#)” on page 9.

### ▼ To associate a Data Change Map to a Data Volume

The `vxassist` command enables you to associate a DCM to a new data volume or an existing data volume.

- ❖ Create the data volume and associate the DCM as follows:

```
# vxassist -g diskgroup make dv_name..... logtype=dcm
```

OR

- ❖ Associate the DCM with an existing data volume as follows:

```
# vxassist -g diskgroup addlog dv_name logtype=dcm
```

VVR mirrors the DCM by default. If `loglen` is not specified, `vxassist` calculates a suitable size for the DCM. For more information, see “[Determining the Region Size](#)” on page 130.

---

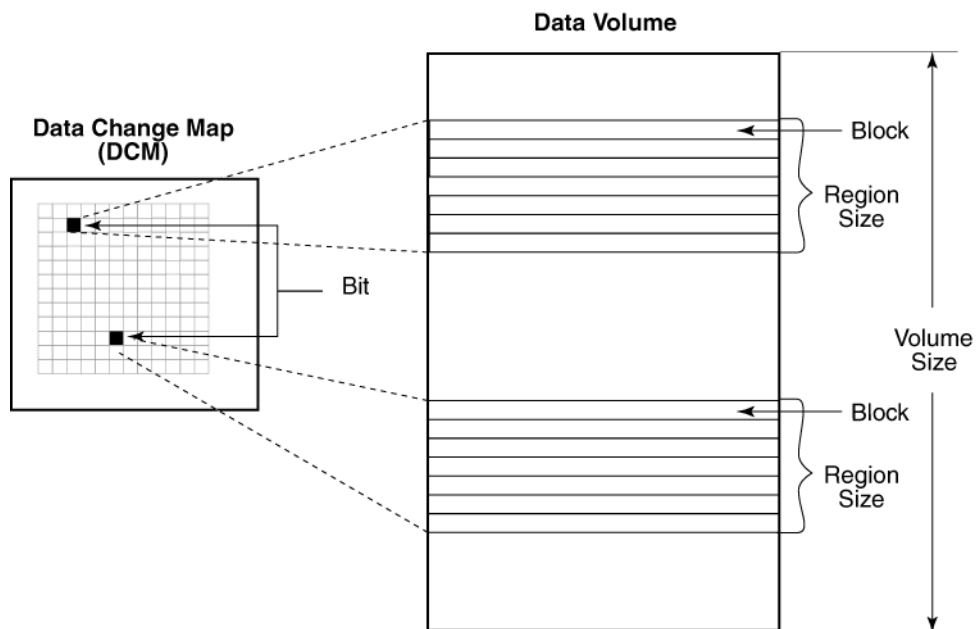
**Note** If you try to grow a volume that has a DCM, an error message warns you if the DCM is not large enough for the increased size. In this case, dissociate the DCM, grow the volume, and then associate a new DCM to the volume.

---



## Determining the Region Size

VVR calculates the DCM size based on the size of the volume. The default size of the DCM ranges from 4K to 256K depending on the size of the volume. However, you can specify the size of the DCM to a maximum of 2 MB. Internally, the DCM is divided into two maps: the active map and the replay map. Each bit in the DCM represents a contiguous number of blocks in the volume it is associated with, which is referred to as the region.



The region size is calculated based on the volume size divided by half the DCM size in bits. The minimum region size is 32 blocks or 32 kilobytes.

The following table gives examples of the region sizes for volumes of different sizes in a non-CDS (Cross-Platform Data Sharing) disk group and a CDS disk group with the default DCM size and a user-specified DCM size of 2 MB.

Volume Size	Non-CDS Disk Group			CDS Disk Group		
	Default DCM Size	Region Size for Default DCM Size	Region Size for a DCM Size of 2 MB Specified by the User	Default DCM Size	Region Size for Default DCM Size	Region Size for a DCM Size of 2 MB Specified by the User.
1 MB	1K	32K	32K	16K	32K	32K
100 MB	1K	32K	32K	16K	32K	32K
200 MB	2K	32K	32K	16K	32K	32K
400 MB	4K	32K	32K	16K	32K	32K
1 GB	9K	32K	32K	16K	32K	32K
2 GB	17K	32K	32K	32K	32K	32K
4 GB	33K	32K	32K	48K	32K	32K
8 GB	65K	32K	32K	80K	32K	32K
20 GB	161K	32K	32K	176K	32K	32K
40 GB	161K	64K	32K	176K	64K	32K
100 GB	201K	128K	32K	208K	128K	32K
200 GB	229K	224K	32K	240K	224K	32K
400 GB	247K	416K	64K	256K	416K	64K
1 TB	249K	1056K	160K	256K	1056K	160K



## Resizing a Data Volume in a Replicated Data Set

The `vradmin resizevol` command enables you to resize a data volume in a Replicated Data Set (RDS) even when replication is in progress. The `vradmin resizevol` command resizes the data volumes in all the RVGs in the RDS. The `vradmin resizevol` command can be entered from any host in an RDS.

---

**Caution** To avoid any problems with the file system on the Secondary, run the `vradmin resizevol` command only when the Secondary is up-to-date. VVR replicates changes to the meta data of a file system on the Primary data volumes to the Secondary. If a takeover happens while these changes are yet to be applied to the Secondary data volumes, the size of the file system may not match the size of the underlying data volume and it may not be possible to mount the file system on the new Primary. If this occurs, run the file system-specific commands to recover the file system.

---

### Important Notes on Resizing a Data Volume in a Replicated Data Set

- ◆ If the Primary data volume contains a file system, the `vradmin resizevol` command also resizes the file system using the `vxresize` command. For more information, see the `vxresize (1M)` manual page.
- ◆ The `vradmin resizevol` command pauses replication, resizes the data volume, and then resumes replication.
- ◆ If you want to increase the size of a data volume, make sure there is enough space on the Primary and the Secondary.

---

**Note** When you increase the size of a data volume, the newly added portions on the Primary and Secondary data volumes are not synchronized. In this case, the output of the `vradmin verifydata` command will show that the checksums for the Primary and Secondary data volumes do not match.

---

- ◆ If the `vradmin resizevol` command fails on any of the hosts in the RDS during its execution, the original volume sizes are not restored. This results in volume size mismatch on the Primary and its Secondaries. To correct this mismatch, correct the error condition and then reissue the `vradmin resizevol` command and resume the Secondary RLINKS.

**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.
- ✓ If you want to increase the size of a data volume, make sure there is enough space in the disk group on the Primary and the Secondary by issuing the following command:

```
# vxrdg -g diskgroup free
```

To resize a volume in an RDS:

```
# vradmin -g diskgroup [-f] resizevol local_rvgname volume_name \  
volume_length
```

The argument *local\_rvgname* is the name of the RVG on the local host and represents its RDS. The *-f* option is required if the data volume involved in the *resizevol* operation is being decreased in size.

The argument *volume\_name* is the name of the data volume to be resized.

The argument *volume\_length* is the desired size of the data volume to be resized. You can specify the volume length using the standard length convention. You can specify a prefix of either the plus (+) or minus (-) sign to increase or decrease the data volume size by the specified amount.

**Examples:**

The following examples show how to resize to different lengths an existing volume *hr\_dv01* in all RVGs of the RDS represented by its local RVG *hr\_rvg*. The disk group *hrdg* contains the local RVG *hr\_rvg*.

To resize the volume *hr\_dv01* to 100 gigabytes, type the following command on any host in the RDS:

```
# vradmin -g hrdg resizevol hr_rvg hr_dv01 100G
```

To increase the size of the data volume *hr\_dv01* by 100 megabytes when the Primary and Secondary data volumes are the same size, type the following command on any host in the RDS:

```
# vradmin -g hrdg resizevol hr_rvg hr_dv01 +100M
```

To decrease the size of the data volume *hr\_dv01* by 500 kilobytes (K) when the Primary and Secondary data volumes are the same size, type the following command on any host in the RDS:

```
# vradmin -g hrdg -f resizevol hr_rvg hr_dv01 -500K
```



## Dissociating a Data Volume from its Replicated Data Set

You can remove a data volume from a Replicated Data Set (RDS) using the `vradm delvol` command. The `vradm delvol` command dissociates a data volume from all the RVGs in an RDS; the volumes are not deleted.

The `vradm delvol` command can be entered from any host in an RDS. If the `vradm delvol` command fails on any of the hosts in the RDS during its execution, the original configuration remains unchanged.

### ▼ To remove a data volume from an RDS when the Primary RVG has been stopped

Type the following command on any host in the RDS:

```
# vradm -g diskgroup delvol local_rvgname volume_name
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `volume_name` is the name of the volume to be removed from the RDS.

### ▼ To remove a data volume from an RDS when the Primary RVG has not been stopped

---

**Note** Although you can use `-f` option with `vradm delvol` command to remove a data volume from an RDS when the Primary RVG has not been stopped, this is not the recommended approach. It is recommended that you stop the Primary RVG before proceeding with this command.

---

Type the following command on any host in the RDS:

```
# vradm -g diskgroup -f delvol local_rvgname volume_name
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `volume_name` is the name of the volume to be removed from the RDS.

#### **Example:**

This example shows how to remove a data volume `hr_dv01` from all RVGs of its RDS. The data volume `hr_dv01` resides on the local host `london` on which the command is entered. The data volume `hr_dv01` is associated with the local RVG `hr_rvg`, which belongs to the disk group `hrdg`.

```
# vradm -g hrdg delvol hr_rvg hr_dv01
```

## 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](#)” on page 28.

To enable SRL protection, each data volume in the RDS must have an associated DCM. For more information, section “[Associating the Component Volumes of a Volume Set to an RDS](#)” on page 127. To change the replication setting for SRL protection, see “[Changing the Replication Settings for a Secondary](#)” on page 49.

### Incrementally Synchronizing the Secondary After SRL Overflow

The default protection mode for the SRL is `autodcm` and every 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 command `vradm resync`. Note that you can also use the `cache` or `cachesize` parameters with the `vradm resync` command. Specifying these attributes will cause the command to first create a space-optimized snapshot of the Secondary data volumes before starting the resynchronization. If the RVG on the Secondary has VxVM ISP volumes, then you cannot use the `cachesize` attribute. 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 RLINK is 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 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 volumes before the resynchronization starts by using the following commands:

1. Create the cache object for the data volumes. For more information, refer to the section [“Preparing the RVG Volumes for Snapshot Operation”](#) on page 164. This step is optional if you plan to use the `cachesize` attribute with `vradmin resync` command.
2. To start the resynchronization use the command:

```
# vradmin -g diskgroup [-wait] resync local_rvgname \
    [cache=cacheobj | cachesize=size]
```

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 `vradmin resync` to create one cache object for each snapshot

The parameters `cache` and `cachesize` are optional. If you do not specify either of these parameters then the `vradmin resync` command will resynchronize the Secondary volumes using the DCM replay, without creating the snapshots.

The `-wait` option can be used with the `vradmin resync` command to wait for the synchronization process to complete.

### SRL Overflow Protection With DCM—Flags and Definitions

If the SRL Overflow Protection With DCM feature has been activated, VVR sets the following flag on the corresponding RLINK and its RVG:

Flag Value	Definition
dcm_logging	Log Overflow Protection With DCM has been started and the DCM is in use.



If the `dcm_logging` flag is set on an RLINK or RVG and neither the `resync_started` nor the `resync_paused` flag is set, the resynchronization (`resync`) has not been started. After the `vradmin resync` command has been issued, one or both of the following flags are set:

Flag Value	Definition
<code>resync_started</code>	Resynchronization is in progress, that is, data is being transferred from the Primary to the Secondary.
<code>resync_paused</code>	Resynchronization is paused.

### Prerequisite for incrementally resynchronizing the Secondary

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

#### ▼ To incrementally resynchronize the Secondary

```
# vradmin -g diskgroup resync local_rvgname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

#### Example:

```
# vradmin -g hrdg resync hr_rvg
```

#### ▼ To determine the progress of the incremental synchronization

Determine the progress of the incremental synchronization after SRL overflow by issuing the following command on the Primary host:

```
# vxrlink -g diskgroup status rlink_name
```

The argument `rlink_name` is the name of the Primary RLINK to the Secondary.

The output shows how much data is left to send. To monitor the progress, issue the `vxrlink -i interval status rlink_name` command. For example, to see the status every 5 seconds, issue the following command:

```
# vxrlink -g hrdg -i5 status rlink_name
```

The output resembles:

```
VxVM VVR vxrlink INFO V-5-1-4464 Rlink rlink_name is in AUTOSYNC.
100864 Kbytes remaining.
VxVM VVR vxrlink INFO V-5-1-4464 Rlink rlink_name is in AUTOSYNC.
94464 Kbytes remaining.
VxVM VVR vxrlink INFO V-5-1-4464 Rlink rlink_name is in AUTOSYNC.
76800 Kbytes remaining.
```



## Breaking Off Mirrors Before Incremental Synchronization

During DCM resynchronization, the data volumes on the Secondary are inconsistent and cannot be used to take over the Primary role. To maintain a consistent copy of the data volumes on the Secondary, break off a mirror from each data volume before starting DCM resynchronization. In the case of a disaster, these mirrors can be used to take over the Primary role. If you have FastResync license, make sure FR is set for all the volumes.

When snapshot plexes are available:

To find out if snapshot plexes are available on a data volume, use the `vxprint` command. The output shows the state of the plex as `SNAPDONE`. If a snapshot plex is available for each data volume, use the `vxrvg snapshot` command to take a snapshot of the data volumes in an RVG. If required, the snapshot volumes can be used to take over the Primary role. After the DCM resynchronization is complete, reattach the snapshot plexes back to the original volume using the `vxrvg snapback` command.

When snapshot plexes are not available:

If snapshot plexes are not available, detach mirrors on each of the data volumes on the Secondary using the `vxplex` command. After the DCM resynchronization is complete, reattach the plexes using the `vxplex att` command. To use the data on a detached plex in situations such as takeover, you must create a volume for the detached plex when snapshot plexes are not available.

### Example 1—When snapshot plexes are not available

This example explains how to break off a mirror from a data volume and reattach a plex after the DCM resynchronization is complete. This example uses the volume `hr_dv01` that has two plexes `hr_dv01_01` and `hr_dv01_02`.

1. On the Secondary, detach a plex from the data volume by typing:

```
# vxplex -g hrdg det hr_dv01_02
```

2. When the RLINK reconnects, incrementally synchronize the Secondary by typing:

```
# vradmin -g hrdg resync hr_rvg
```

For multiple Secondary hosts, VVR simultaneously synchronizes all Secondary hosts that are operating in `dcm logging` mode.

3. After the DCM resynchronization is complete, reattach the plex to the data volume on the Secondary by typing:

```
# vxplex -g hrdg att hr_dv01 hr_dv01_02
```

## Example 2—When snapshot plexes are not available and a disaster occurs

If during the resynchronization process, a disaster occurs and the Secondary takes over, you can recreate the volumes as they were before the resynchronization started. The example uses the RVG `hr_rvg` and the volume `hr_dv01` as in “[Example 1—When snapshot plexes are not available](#)” on page 138. All steps are performed on the former Secondary, which is now the Primary (to take over the Primary role, see “[Taking Over from an Original Primary](#)” on page 223).

1. Detach the Secondary RLINK.

```
# vxrlink -g hrdg det rsec
```

2. Dissociate the original data volume from the Secondary RVG.

```
# vxvol -g hrdg dis hr_dv01
```

3. Remove the original data volume.

```
# vxedit -g hrdg -rf rm hr_dv01
```

4. Create the volume for the detached plex by typing:

```
# vxmake -g hrdg -U usetype vol hr_dv01 plex=hr_dv01_02
```

If a volume contains a file system, specify the *usetype* as `fsgen`; otherwise, specify `gen`.

5. Start the data volume by typing:

```
# vxvol -g hrdg -f start hr_dv01
```

6. Associate the data volume to its RVG.

```
# vxvol -g hrdg assoc hr_rvg hr_dv01
```

7. The volume is no longer mirrored. To add mirrors, issue the following command:

```
# vxassist -g hrdg mirror hr_dv01
```



### Notes on using Incremental Synchronization on SRL Overflow

- ◆ Each data volume in the Primary RVG must have a DCM associated with it. You cannot use the SRL Overflow Protection With DCM feature unless every data volume in the RVG has a DCM. If any of the data volumes in the RVG do not have a DCM, you cannot set `srlprot=dc` or `srlprot=autodcm`. An attempt to associate a volume without a DCM to an RVG that has an RLINK with `srlprot=dc` or `srlprot=autodcm` will also fail.
- ◆ If an RLINK is undergoing Automatic Synchronization and an attached RLINK with SRL Overflow Protection With DCM is about to overflow, the Automatic Synchronization is abandoned and SRL Overflow Protection With DCM for the overflowing RLINK becomes active.
- ◆ If an existing RLINK is using the DCM mechanism and another existing RLINK is about to overflow, the second RLINK is detached unless the DCM resynchronization for the first RLINK has not yet sent any writes. In this case, the outstanding writes of the first RLINK are also sent on the second RLINK.
- ◆ To remove a Secondary from a DCM resynchronization process, detach the corresponding Primary RLINK.
- ◆ If you try to dissociate a DCM from a data volume while the DCM is in use, the operation fails.
- ◆ If the DCM is detached because of I/O errors while the DCM is in use, the resynchronization is abandoned and all the RLINKs that are being synchronized are detached.

## Resizing the SRL

The size of the SRL must be large enough to meet the constraints explained in the section “Sizing the SRL” in the *VERITAS Volume Replicator Planning and Tuning Guide*. These constraints can change with the changes in the business needs, application write rate, available network bandwidth, and so on. As a result, it is necessary to determine the appropriate size of the SRL again. This section includes the following tasks:

- ✓ [Increasing the Size of the SRL on the Primary](#)
- ✓ [Decreasing the Size of the SRL on the Primary](#)
- ✓ [Resizing the SRL on the Secondary](#)

### Increasing the Size of the SRL on the Primary

VVR enables you to increase the size of the Primary SRL in a Replicated Data Set (RDS) using the `vradm resize` command, even while the application is active or while replication is in progress. The `vradm resize` command does not increase the size of the Secondary SRL.

Before increasing the size of the Primary SRL, do the following:

- ✓ Check whether there is enough free space in the disk group in which the SRL resides by issuing the following command:

```
# vxdbg -g diskgroup free
```

#### ▼ To increase the size of the SRL on the Primary

Issue the following command on the Primary:

```
# vradm -g diskgroup resize local_rvgname length
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `length` is the desired size for the Primary SRL. The length can be specified using the standard VxVM conventions. It can be prefixed by a plus sign (+) to indicate an increase in the size of the Primary SRL by the specified amount.

### Decreasing the Size of the SRL on the Primary

For instructions on decreasing the size of the SRL on the Primary, see “[Decreasing the Size of the SRL on the Primary](#)” on page 404.



## Resizing the SRL on the Secondary

Note that you do not have to stop the applications while resizing the Secondary SRL.

### ▼ To resize the SRL on the Secondary

1. Detach the RLINK:

```
# vxrlink -g diskgroup det rlink_name
```

2. Dissociate the SRL from the RVG.

---

**Note** Any checkpoints that you have created will be lost after dissociating the SRL.

---

```
# vxvol -g diskgroup dis srl_name
```

3. Resize the SRL using the `vxassist` command. For example, to increase the size of the SRL:

```
# vxassist -g diskgroup growto srl_name new_length
```

---

**Note** We recommend that the SRL reside on disks that are not being used for the data volumes. We also recommend that the Primary and the Secondary SRLs be the same size. See the `vxassist (1M)` manual page for more information.

---

4. Reassociate the SRL with the RVG:

```
# vxvol -g diskgroup aslog rvg_name srl_name
```

5. Attach the RLINK:

```
# vxrlink -g diskgroup att rlink_name
```

# Administering Replication

## Changing the Replication Settings

You can change the VVR replication attributes according to your requirements with the `vradmin set` command. The `vradmin set` command enables you to set the following VVR replication attributes:

- ◆ Replication Mode
- ◆ Latency protection
- ◆ SRL protection
- ◆ Network transport protocol
- ◆ Packet size
- ◆ Bandwidth limit

For more information on replication settings, see [“Changing the Replication Settings for a Secondary”](#) on page 49.

## Pausing and Resuming Replication to a Secondary

Pausing an RLINK prevents new and already-queued updates from reaching the Secondary from the Primary, and the Primary and Secondary do not communicate.

The `vradmin pauserrep` command does not provide a way to pause a Secondary RLINK. To do this, use the `vxrlink` command on the Secondary host. The `vradmin resumerep` command resumes both types of pauses on the selected RLINKs.

---

**Note** If the latency protection is set to `override`, read [“Primary and Secondary Disconnected”](#) on page 32.

---



### ▼ To pause and resume replication to a Secondary

To pause replication to a specific Secondary in an RDS, specify the Secondary hostname. For an RDS with a single Secondary, you do not have to specify the Secondary hostname.

On any host in the RDS:

1. Pause replication.

```
# vradmin -g diskgroup pauserep local_rvgname [sec_hostname]
```

where *local\_rvgname* is the name of the RVG on the host where the command is issued, and *sec\_hostname* is the name of the host to which replication is being paused.

2. On the Primary, issue the `vxprint` command to check that the state of the RLINK is PAUSE.

```
# vxprint rlink_name
```

3. Resume replication to the Secondary.

```
# vradmin -g diskgroup resumerep local_rvgname [sec_hostname]
```

where *local\_rvgname* is the name of the RVG on the host where the command is issued, and *sec\_hostname* is the name of the host to which replication is being resumed.

## Stopping Replication to a Secondary

The `vradmin stoprep` command can be used to stop replication to a Secondary in an RDS. The `vradmin stoprep` command can be entered from any host in the RDS.

The `vradmin stoprep` command fails if the Primary and Secondary RLINKs are not up-to-date. Use the `-f` option to stop replication to a Secondary even when the RLINKs are not up-to-date.

Before stopping replication, the `vradmin stoprep` command displays a warning and prompts the user to confirm whether or not to stop replication. To skip this confirmation, use the `-s` option with the `vradmin stoprep` command. The `-s` option to the `vradmin stoprep` command proves useful in scripts.

To stop replication to a specific Secondary in an RDS, specify the Secondary host name. For an RDS with a single Secondary, you do not have to specify the Secondary host name.



### ▼ To stop replication to a specific Secondary in an RDS

```
# vradmin -g diskgroup stoprep local_rvgname sec_hostname
```

The argument *local\_rvgname* is the name of the RVG on the local host and represents its RDS.

The argument *sec\_hostname* is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

#### Example:

To stop replication from the Primary RVG `hr_rvg` on `seattle` to the Secondary RVG on host `london`, type:

```
# vradmin -g hrdg stoprep hr_rvg london
```

## Changing the IP Addresses Used for Replication

You may need to change the host name or IP address of the Primary and Secondary used for replication if you move a Primary or Secondary to a new location or if you need to make the replication use a different network. You can change the host name or IP address even after replication has been established. The `vradmin changeip` command enables you to change the replication network between the Primary and a Secondary in an RDS.

#### Prerequisites for changing the IP addresses used for replication:

- ✓ The new host names must be configured for proper resolution at both the Primary and Secondary sites using the appropriate mechanisms such as DNS, NIS, or hosts. This means that each system must be configured to bring up their addresses on reboot, or if this is a cluster, the cluster monitor must bring up the proper address.
- ✓ 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 `vradmin changeip` command must be run from the Primary host.

---

**Note** The VVR heartbeat port can be changed using the `vrport` command. To ensure that the RLINKs pick up the new port, always run the `vradmin changeip` command (without passing the *newpri* and *newsec* arguments) after changing the port. Restart the `vxnetd` daemon on the required system for the changes to take effect.

---



### ▼ To change the IP addresses used for replication

```
vradmin [-g diskgroup] changeip local_rvgname [sec_hostname] \  
[newpri=<new_pri_ip | hostname>] [newsec=<new_sec_ip | hostname>]
```

The argument *diskgroup* is the name of the local disk group that contains the RVG.

The *local\_rvgname* argument is the name of the RVG on the host where the command is issued.

The *sec\_hostname* is the name of the Secondary to which the replication network is being changed. This argument must be specified if the RDS has more than one Secondary.

The *newpri* attribute specifies a new hostname or IP address for the Primary host that is to be used to establish a network connection for the replication to the Secondary. This is the new value for the *local\_host* attribute of the Primary RLINK and the *remote\_host* attribute of the corresponding Secondary RLINK.

The *newsec* specifies a new hostname or IP address for the Secondary host that is to be used to establish a network connection for the replication. This is the new value for the *remote\_host* attribute of the Primary RLINK and the *local\_host* attribute of the Secondary RLINK.

#### Example:

This example shows how to change the network used for replication. The current configuration looks like this:

	On Primary	On Secondary
<b>local_host</b> displayed in the output of the <code>vxprint -l <i>rlink_name</i></code> command	seattle	london
<b>remote_host</b>	london	seattle
<b>RVG</b>	hr_rvg	hr_rvg
<b>Disk Group</b>	hrdg	hrdg
<b>RLINK</b>	rlk_london_hr_rvg	rlk_seattle_hr_rvg

After making the changes to the replication network, the configuration looks like this:

	On Primary	On Secondary
<b>local_host</b> displayed in the output of the <code>vxprint -l rlink_name</code> command	seattle_hrnet	london_hrnet
<b>remote_host</b>	london_hrnet	seattle_hrnet
<b>RVG</b>	hr_rvg	hr_rvg
<b>Disk Group</b>	hrdg	hrdg
<b>RLINK</b>	rlk_london_hr_rvg	rlk_seattle_hr_rvg

### ▼ To change the IP addresses used for replication

1. From the Primary host `seattle`, issue the following command:

```
# vradmin -g hrdg changeip hr_rvg newpri=seattle_hrnet \  
newsec=london_hrnet
```

The `vradmin changeip` command changes the IP address of both the Primary RLINK and the corresponding Secondary RLINK to the new addresses `newpri` and `newsec` (in this example, `seattle_hrnet` and `london_hrnet`, respectively).

2. To verify the change on the Primary RLINK, issue the following command on the Primary host:

```
# vxprint -l rlk_london_hr_rvg
```

Output includes the following fields:

```
Disk group: hrdg  
.  
.  
Rlink: rlk_london_hr_rvg  
.  
.  
remote_host=london_hrnet IP_addr=x.x.x.x  
.  
.  
local_host=seattle_hrnet IP_addr=x.x.x.x  
.  
.
```

where `x.x.x.x` represents the corresponding IP address.



- To verify the change on the Secondary RLINK, issue the following command on the Secondary host:

```
# vxprint -l rlk_seattle_hr_rvg
```

Output includes the following fields:

```
Disk group: hrdg
.
.
Rlink: rlk_seattle_hr_rvg
.
.
remote_host=seattle_hrnet IP_addr=x.x.x.x
.
.
local_host=london_hrnet IP_addr=x.x.x.x
.
.
```

where x.x.x.x represents the corresponding IP address.

## Changing the Network Ports Used for Replication

### Port Numbers Used by VVR

VVR uses the UDP and TCP transport protocols to communicate between the Primary and Secondary. This section lists the default ports used by VVR. VVR uses the following ports by default when replicating data using UDP:

Port Numbers	Description
UDP 4145	IANA approved port for heartbeat communication between the Primary and Secondary.
TCP 8199	IANA approved port for communication between the <code>vradmin</code> daemons on the Primary and the Secondary.
TCP 8989	Communication between the <code>in.vxrsyncd</code> daemons, which are used for differences-based synchronization.
UDP Anonymous ports (OS dependent)	Ports used by each RLINK for data replication between the Primary and the Secondary.

VVR uses the following ports by default when replicating data using TCP.

Port Numbers	Description
UDP 4145	IANA approved port for heartbeat communication between the Primary and Secondary.
TCP 4145	IANA approved port for TCP Listener port.
TCP 8199	IANA approved port for communication between the <code>vradmin</code> daemons on the Primary and the Secondary.
TCP 8989	Communication between the <code>in.vxrsyncd</code> daemons, which are used for differences-based synchronization.
TCP Anonymous ports	Ports used by each RLINK for replication on the Primary.

## Displaying and Changing the Ports Used by VVR

Use the `vrport (1M)` command to display, change or set the port numbers used by VVR. You may have to change the port numbers in the following cases:

- ✓ To resolve a port number conflict with other applications.
- ✓ To configure VVR to work in your firewall environment.
- ✓ To configure VVR to work in your firewall environment when using UDP; to specify a restricted number of ports to replicate data between the Primary and the Secondary.

### Port Used for Heartbeats

Use the `vrport heartbeat` command to display the port number used by VVR, for heartbeats. To change the heartbeat port number on a host, specify the port number with the `vrport heartbeat` command. Use the `vradmin changeip` command to update the RLINKs with the new port information, and then restart the `vxnetd` daemon on the required system for the changes to take effect.

#### ▼ To display the port number used for heartbeats

```
# vrport heartbeat
```

#### ▼ To change the port number used for heartbeats

```
# vrport heartbeat port
```



## Example

### ▼ To change the replication heartbeat port on seattle from 4145 to 5000

1. Use the `vrport` command to change the heartbeat port to 5000 on the required host.

```
# vrport heartbeat 5000
```

2. Issue the `vradmin changeip` command without the `newpri` and `newsec` attributes.

```
# vradmin -g hrdg changeip hr_rvg london
```

3. Verify the changes to the local RLINK by issuing the following command on the required host:

```
# vxprint -g hrdg -l rlk_london_hr_rvg
```

4. Stop the `vxnetd` daemon.

```
# /usr/sbin/vxnetd stop
```

5. Restart the `vxnetd` daemon.

```
# /usr/sbin/vxnetd
```

Follow the same steps to change the heartbeat port on secondary (london).

---

**Note** VVR supports a configuration with different heartbeat port numbers on the primary and secondary.

---

### Port Used by `vradmin`

To display the port number used by `vradmin`, use the `vrport vradmin` command. To change the `vradmin` port, specify the port number with the `vrport vradmin` command.

### ▼ To display the port number used by `vradmin`

```
# vrport vradmin
```

### ▼ To change the port number used by `vradmin`

```
# vrport vradmin port
```

---

**Note** You must restart the server `vradmin` for this change to take effect. Make sure you change the port number on all the hosts in the RDS.

---

### Port Used by `in.vxrsyncd`

To display the port numbers used by `in.vxrsyncd`, use the `vrport vxrsyncd` command. To change the port numbers used by `in.vxrsyncd`, specify the port number with the `vrport vxrsyncd` command.

#### ▼ To display the port number used by `in.vxrsyncd`

```
# vrport vxrsyncd
```

#### ▼ To change the port number used by `in.vxrsyncd`

```
# vrport vxrsyncd port
```

---

**Note** You must restart the server `in.vxrsyncd` for this change to take effect. Make sure you change the port number on all the hosts in the RDS.

---

### Ports Used to Replicate Data Using UDP

To display the ports used to replicate data when using UDP, use the `vrport data` command. To change the ports used to replicate data when using UDP, specify the list of port numbers to use with the `vrport data` command.

Each RLINK requires one UDP port for replication. Specify an unused, reserved port number that is less than 32768 so that there is no port conflict with other applications. The number of ports specified must be equal to or greater than the number of RLINKs on the system.

---

**Note** For systems using the TCP protocol for replication, you are not required to select any data port as the connection is established with the listener port on the remote host. The listener uses this port number which is numerically same as the UDP port used for heartbeat messages.

---

#### ▼ To display ports used to replicate data when using UDP

```
# vrport data
```

#### ▼ To change ports used to replicate data when using UDP

For a system configured with one RLINK, use the following command:

```
# vrport data port
```

For a system configured with multiple RLINKs, you can specify either a range of port numbers or a list of port numbers or both.

To specify a range of port numbers, use the following command:

```
# vrport data port1, port2, portlow-porthigh, .....
```



For example:

```
# vreport data 3400, 3405, 3500-3503, 5756-5760
```

---

**Note** To use the new port information, execute `/usr/sbin/vxnetd`, and then pause and resume all RLINKs.

---

## Administering the Replicated Data Set

### Removing a Secondary from a Replicated Data Set

The `vradm delsec` command removes a Secondary RVG from its RDS. The `vradm delsec` command can be entered from any host in an RDS.

The `vradm delsec` command removes the Secondary RVG from its RDS on the specified Secondary host. Before executing this command, you must stop replication to the specified Secondary, using the `vradm stoprep` command.

---

**Caution** The operation performed by the `vradm delsec` command is irreversible.

---

The `vradm delsec` command performs the following by default:

- ◆ Dissociates the data volumes and SRL from the Secondary RVG.
- ◆ Removes the Secondary RVG from its RDS, deletes the Secondary RVG, and deletes the associated Primary and Secondary RLINKs.

The `vradm delsec` command does not delete data volumes and the SRL.

#### ▼ To remove a Secondary from an RDS

```
# vradm -g diskgroup delsec local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradm printrvg` command.



**Example:**

This example removes the Secondary RVG `hr_rvg` from its RDS when the configuration looks like this:

	On Primary	On Secondary
<b>Host Name</b> displayed in the output of the <code>vradm</code> <code>printrvg</code> command	seattle	london
<b>RVG</b>	hr_rvg	hr_rvg
<b>Disk Group</b>	hrdg	hrdg

Because the command is entered on the Primary `seattle`, the local RVG is the Primary RVG `hr_rvg` belonging to the disk group `hrdg`.

To remove the Secondary RVG `hr_rvg` from its RDS, type the following command on `seattle`:

```
# vradm -g hrdg delsec hr_rvg london
```

## Removing a Primary RVG

The `vradm delpri` command removes a Primary RVG from an RDS and thus deletes 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 152.

The `vradm delpri` command performs the following by default:

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

The `vradm delpri` command does not delete data volumes or the SRL from the VERITAS Volume Manager configuration.

---

**Note** This command can only be issued from the Primary host.

---

### ▼ To remove a Primary RVG

```
# vradm -g diskgroup delpri rvg_name
```

The argument `rvg_name` is the name of the Primary RVG to be removed.



When used with the `-f` option, the `vradm delpri` command removes the Primary RVG even when the application is running on the Primary.

**Example 1:**

To remove the Primary RVG `hr_rvg` when the application using the Primary data volumes is inactive, issue the following command on the Primary host:

```
# vradm -g hrdg delpri hr_rvg
```

**Example 2:**

To remove the Primary RVG `hr_rvg` when the application using the Primary data volumes is active, issue the following command on the Primary host:

```
# vradm -g hrdg -f delpri hr_rvg
```

## Administering Checkpoints

A checkpoint is a user-defined marker in the SRL that can be used to synchronize the Secondary when the application is active, or to restore the Secondary data volumes from backup. For information on using checkpoints for synchronization, see [“Synchronizing the Secondary and Starting Replication”](#) on page 58. For information on using checkpoints for backing up and restoring a Secondary, see [“Backing Up the Secondary”](#) on page 191.

This section describes how you can work with checkpoints. For general information about checkpoints, see [“Understanding Checkpoints”](#) on page 25.

## Creating Checkpoints

VVR enables you to create Primary and Secondary checkpoints. The Primary checkpoints are associated with an RVG. However, Secondary checkpoints are associated with an RLINK. VVR allows you to create a maximum of 46 checkpoints.

**▼ To create a Primary checkpoint**

```
# vxrvrg -c checkpt_name checkstart rvg_name
```

The argument `checkpt_name` is the name that you choose to specify for the checkpoint.

**▼ To create a Secondary checkpoint**

```
# vxrlink -c checkpt_name pause rlink_name
```

The argument `checkpt_name` is the name that you choose to specify for the checkpoint.

## Ending Checkpoints

The end of the checkpoint, or checkend, marks the position in the SRL for the end of a process such as synchronization or backup. When you are ready to end the checkpoint, for example, when the backup is complete, end the checkpoint in the SRL.

### ▼ To end a Primary checkpoint

```
# vxrvg -g diskgroup checkend rvg_name
```

### ▼ To end a Secondary checkpoint

For a Secondary checkpoint, resume replication to the Primary.

```
# vxrlink -c checkpt_name resume rlink_name
```

The checkpoint will end when the resume operation starts. For more information about using Secondary checkpoints, see [“Backing Up the Secondary”](#) on page 191.

## Viewing Checkpoints

Primary checkpoints are associated with an RVG. You can display the list of Primary checkpoints by using the `vxrvg cplist` command.

Secondary checkpoints on the other hand are associated with an RLINK. You can display the list of Secondary checkpoints by using the `vxrlink cplist` command on the Primary.

For more information on these commands, refer to [“Displaying a List of Checkpoints”](#) on page 91.

## Deleting Checkpoints

After you have finished using the checkpoints, you can delete the checkpoints. VVR allows you to retain a maximum of 46 checkpoints. To create any new checkpoint, delete an earlier checkpoint that you no longer require.

### ▼ To delete a Primary checkpoint

```
# vxrvg -g diskgroup -c checkpt_name checkdelete rvg_name
```

The argument `rvg_name` is the name of the Primary RVG for which the checkpoint is to be deleted.

The argument `checkpt_name` is the name of the specific checkpoint to be deleted.



▼ **To delete a Secondary checkpoint**

```
# vxrlink -g diskgroup -c checkpt_name checkdelete rlink_name
```

Note that this command must be run only on the Primary.

The argument *rlink\_name* is the name of the RLINK for which the checkpoint is to be deleted.

The argument *checkpt\_name* is the name of the specific checkpoint to be deleted.

## Creating RVG Snapshots

VVR enables you to create snapshots that are images of the online 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. VVR provides two methods of creating snapshots: instant snapshots 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 is a separately licensed feature of VVR. The advantages of this method over the traditional snapshot method are that the snapshots are available immediately, and they may be space-optimized, thus requiring less space than the traditional snapshots. For more information on instant snapshots, see [“Using the Instant Snapshot Feature”](#) on page 157.

With the traditional snapshot method, depending on the size of the volume, the time required for initial synchronization of the plexes can be very large. For more information on traditional snapshots, see [“Using the Traditional Snapshot Feature”](#) on page 177.

After a volume has been prepared for the instant snapshot feature, you cannot use it to take a snapshot using the traditional snapshot method. To use the traditional snapshot method, you must first unprepare the volume. Thus, you cannot use the same volume for both the traditional method and the instant snapshot method at the same time.

If an RVG contains component volumes of a volume set, the `vrxvg snapshot` command cannot be used to take snapshots of its data volumes. To take snapshots of the data volumes in an RVG that contains component volumes of a volume set, use the `vxsnap make` command.

If an RVG consists of VxVM ISP volumes, then you can create only instant full or space-optimized snapshots of the volumes in the RVG. When creating instant space-optimized snapshots, make sure you create the cache object beforehand because the `cachesize` parameter is not supported for an RVG with VxVM ISP volumes.

## Using the Instant Snapshot Feature

VVR enables you to create instant snapshots using the `vrxvg snapshot` command. This command 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. When creating full instant or space-optimized snapshots, the snapshot volumes do not need to be synchronized beforehand, therefore the snapshots are available instantly. The snapshot volumes are then synchronized later in the background.



## Snapshot Naming Conventions

The snapshot volumes must be created using the correct naming convention, that is, `<prefix>-dv_name`. You can create snapshot volumes with appropriate prefixes using the `-P` option. However, you must ensure that this prefix matches the prefix that you specified for the data volumes.

For example, if you specify the prefix as `month`, the name of each snapshot data volume will start with the prefix `month`; that is, it will be named as `month-dv_name`. Thus, a data volume `hr_dv01` can have snapshot volumes such as `june-hr_dv01`, `july-hr_dv01`.

---

**Note** We recommend that you create snapshots with prefixes using the `-P` option so that the snapshots can be easily identified for restoring the data. However, if you do not specify any prefix, the default prefix `SNAP` will be used.

---

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

- ◆ [Instant Full Snapshot](#)
- ◆ [Instant Space-Optimized Snapshots](#)
- ◆ [Instant Plex-breakoff Snapshots](#)

---

**Note** If the RVG consists of VxVM ISP volumes then VVR supports only the instant full snapshot and instant space-optimized snapshot operations.

---

## Instant Full Snapshot

The `vrxvg -F snapshot` command enables you to create an instant full snapshot of all the volumes in the RVG, at a single point in time. The snapshot is available for use immediately, because the snapshot volumes do not have to be completely synchronized, in the beginning. The snapshots volumes are synchronized later in the background.

The `vrxvg snapshot` command creates the data volume snapshots for all the volumes in the RVG, similar to the ones created by the `vxsnap make` command.

---

**Note** You must create and prepare the snapshot volumes before they can be used for creating snapshots.

---

### Prerequisites for creating instant full snapshots

- ✓ Make sure you create the snapshot volumes and then prepare them before creating the snapshots.
- ✓ Make sure the snapshot volumes are the same size as the original volumes.
- ✓ Make sure the snapshot volumes follow a proper naming convention such that the snapshot volume names can be easily linked to the original volumes. For more information, refer to the section “[Snapshot Naming Conventions](#)” on page 158.

The steps required to create instant full snapshots are as follows:

- ◆ Creating Snapshot Volumes for Data Volumes in an RVG
- ◆ Preparing the Volume
- ◆ Freezing or Pausing Replication
- ◆ Taking a Snapshot
- ◆ Unfreezing or Resuming Replication

### Creating Snapshot Volumes for Data Volumes in an RVG

You must create snapshot volumes for the data volumes in an RVG before you take an instant full snapshot, because the `vrxvg snapshot` command does not create the snapshot volumes. Use the `vxassist make` commands or other Volume Manager commands to create the required volumes. For more information on creating the volumes, refer to the *VERITAS Volume Manager Administrator's Guide*.

### Preparing the Volumes

#### Upgrading the disk group

To use the instant snapshot feature, the disk groups must be version 120 or above. If you are using disk groups created with an earlier disk group version, the first step in preparing the volumes is to upgrade the disk groups.

To explicitly upgrade the disk groups, run the `vxdg upgrade` command. For details on upgrading disk groups, refer to the *VERITAS Volume Manager Administrator's Guide*.



### Preparing volumes for an instant snapshot

Use the following command to prepare the volumes for instant snapshots:

```
# vxsnap -g diskgroup prepare volume [region=size] \  
[ndcomirs=number] [storage_attribute...]
```

---

**Note** Run this command once on every data volume in the RVG.

---

After running the command the volumes are ready for the instant snapshot operations. Use the following command to verify whether the volumes are prepared. This command checks the settings of the *instant* flag for the required volumes:

```
# vxprint -g diskgroup -F%instant <volume>
```

For more information on the `vxsnap prepare` command, refer to the *VERITAS Volume Manager Administrator's Guide*.

### Freezing or Pausing Replication

Before taking the snapshot on the Secondary, make sure the data volumes are consistent at the application level by either freezing or pausing replication. To make the data volumes consistent at the application level, use the IBC messaging utility `vxibc`. For information on using the `vxibc` utility, refer to [“Using VVR for Off-Host Processing”](#) on page 199. For a failed Primary, you can pause the Primary RLINK, and then take a snapshot of the RVG. If you do not use `vxibc`, you pause the RLINK before taking the snapshot.

VVR provides you with sample scripts that can be used to freeze the replication before creating instant snapshots. When you install VVR, these scripts are installed in the following directory:

```
/etc/vx/vvr/ibc_scripts/sample_so_snapshot
```

Refer to the *README* file in this directory for instructions on how to use the sample scripts to create the instant snapshots.



## Creating Instant Full Snapshots

Use the following command to create an instant full snapshot for each data volume in an RVG:

```
# vxrvrg -g diskgroup [-P prefix] -F snapshot rvg_name \  
  [instantso=volume_list {cache=cachename|cachesize=size}] \  
  [plexbreakoff=volume_list [plexprefix=plex_prefix]] \  
  [exclude=volume_list] [syncing=yes|no] [comment="<comment>"]
```

Use the `vxrvrg snapshot` command with its different attributes to specify the type of snapshot that you want to create. The `-F` option specifies instant full snapshots. By default, all the volumes in the RVG are considered for the instant full snapshots. To exclude any of the volumes from being considered for the instant full snapshots, use one of the following attributes. Depending on the attribute specified with the `vxrvrg snapshot` command, the appropriate snapshots of the volumes are created.

The attribute `instantfull` need not be specified when the `-F` option is specified. This option can be used only if the `-F` option has not been specified and you still require a full instant snapshot of some of the volumes.

The attribute `exclude` specifies a comma-separated list of volumes that do not need to be considered for any kind of snapshots.

The attribute `instantso` specifies a comma-separated list of volumes that can be included for instant space-optimized snapshots when taking the instant full snapshots.

The attribute `syncing` specifies whether you want to start synchronizing the volumes in the background. By default, the value for this attribute is `yes`. The synchronization process can be started or stopped as required. For more information, refer to [“Synchronizing Volumes on the Local Host and Remote Hosts”](#) on page 124.

By default, when using the `vxrvrg snapshot` command with the `-F` option you need not specify the volume list since all the volumes are included. However, you can also specify the volume list for the `instantso`, `plexbreakoff` and the `exclude` attribute when using the `-F` option. This will result in some volumes having an instant full snapshot, some having an instant space-optimized snapshot and some of them being excluded.

Any volumes in the RVG that are not specified in the `volume_lists` of the attributes `exclude`, `plexbreakoff`, or `instantso` will be snapped in the same way as the specified snapshot type, which is the instant full snapshot. Note that it is not possible to create two different types of snapshots for the same volume. The snapshot operation may result in some volumes with no snapshots and some with one of the three types of snapshots.

You can also reattach the snapshots to the data volumes in the RVG using the `vxrvrg snapback` command. For more information on reattaching the snapshots, refer to [“Reattaching the Snapshot Plexes to the Data Volumes \(snapback\)”](#) on page 171.



**Example:**

To specify a prefix for the snapshots of each data volume in an RVG, use the following command. Make sure the snapshot volumes have been created in advance.

```
# vxrvrg -g diskgroup -P june -F snapshot rvg_name
```

A snapshot data volume with the name *june-dv\_name* is created for each data volume in the RVG. You can have more than one snapshot of the data volumes in an RVG.

**Example:**

This example describes the steps to create an instant full snapshot for an RVG:

1. Find the size of the original volumes for which you want to create snapshots on the hosts `seattle` or `london`.

```
# vxprint -g hrdg -F"%name %len" hr_dv01 hr_dv02
```

OR

```
# vxprint -g hrdg -F"%name %len" `vxrvrg getdatavols hr_rvg`
```

2. Prepare the volumes for the instant snapshot operations, by using the following command on each data volume in the RVG for which you intend to create snapshots:

```
# vxsnap -g hrdg prepare hr_dv01
```

```
# vxsnap -g hrdg prepare hr_dv02
```

---

**Note** Make sure that all the applications running on these volumes are closed.

---

3. Create the snapshot volumes of the same size (or greater) as the original volumes and with an appropriate prefix.

```
# vxassist -g hrdg make JUNE-hr_dv01 20971520
```

```
# vxassist -g hrdg make JUNE-hr_dv02 20971520
```

4. Prepare the snapshot volumes:

```
# vxsnap -g hrdg prepare JUNE-hr_dv01
```

```
# vxsnap -g hrdg prepare JUNE-hr_dv02
```

5. Pause or freeze replication on the Secondary.

For more details on pausing replication, refer to [“Pausing and Resuming Replication to a Secondary”](#) on page 143. For more information on freezing replication, refer to [“Using the In-Band Control Messaging Utility vxibc and the IBC Programming API”](#) on page 371.

6. Create a snapshot of the required volumes in the RVG:

```
# vxrvg -g hrdg -F -P JUNE snapshot hr_rvg
```

This command creates a snapshot of all the volumes in the RVG with a prefix JUNE.

7. Resume or unfreeze replication on the Secondary.

For more information on resuming replication, refer to [“Pausing and Resuming Replication to a Secondary”](#) on page 143. For more information on unfreezing replication, refer to [“Using the In-Band Control Messaging Utility vxibc and the IBC Programming API”](#) on page 371.

## Unfreezing or Resuming Replication

After taking a snapshot, unfreeze replication if you are using IBC messaging; otherwise, if you have paused replication, resume it. The snapshots are now ready for use.

## Instant Space-Optimized Snapshots

The `vxrvg -S snapshot` command creates an instant space-optimized snapshot of all the volumes in the RVG at a single point in time. The `vxrvg snapshot` command creates the same type of snapshots as the `vxsnap make` command and uses a cache object that functions as a space-optimized persistent store. The space required by the space-optimized snapshots is less than that of the original volume because space-optimized snapshots store only the changed data. The data change between the source volume and the snapshot usually is minimal during the lifetime of the snapshot.

If the size of the cache object is not enough for the incoming writes, the cache object can grow in size automatically, provided that the `autogrow` attribute has been set to `on`.

The values of the `highwatermark`, `autogrowby` and `maxautogrow` attributes can be set when a cache object is created using `vxmake`. If necessary, you can use the `vxcache set` command to change the values of these attributes for an existing cache. The default value for these attributes is as follows:

Attribute	Default value
<code>autogrow</code>	Default is <code>off</code> .
<code>autogrowby</code>	Default value is 20% of the size of the cache volume in blocks.
<code>highwatermark</code>	Default is 90% of the size of the cache volume in blocks.
<code>maxautogrow</code>	Default is twice the size of the cache volume in blocks.



When the cache volume that is used by the snapshot reaches the preset *highwatermark* value, the VERITAS Volume Manager cache daemon, *vxcached*, is invoked. The values of the *highwatermark*, *autogrowby* and *maxautogrow* attributes for the cache object determine the behavior of *vxcached* daemon.

- ◆ If the cache usage reaches the *highwatermark* value and the new required cache size cannot exceed the value of *maxautogrow* then *vxcached* grows the size of the cache volume by the size *autogrowby*.
- ◆ When cache usage reaches the *highwatermark* value, and the value of the new cache that needs to be created exceeds the value of *maxautogrow*, then *vxcached* deletes the oldest snapshot in the cache. If there are several snapshots that have been created during the same time period, then the largest of these is deleted.
- ◆ If the autogrow feature has been disabled for the cache object and the cache usage reaches the *highwatermark* value, then *vxcached* deletes the oldest snapshot in the cache. If there are several snapshots that have been created during the same time period, the largest of these is deleted. If there is only a single snapshot, the snapshot is detached and marked as invalid.

For more information on the *vxcached* daemon or the attributes of the autogrow parameter, refer to the *VERITAS Volume Manager Administrator's Guide*.

The *vrxvg snapshot* command also enables you to specify the size for the cache using the *cachesize* parameter. In this case, a separate cache object is created for every space-optimized snapshot.

---

**Note** If the RVG consists of only VxVM ISP volumes then the *cachesize* parameter cannot be used.

---

The steps required to create space-optimized snapshots are as follows:

- ◆ Preparing the RVG volumes for snapshot operation
- ◆ Creating the cache object
- ◆ Freezing or Pausing Replication
- ◆ Taking a space-optimized snapshot
- ◆ Unfreezing or Resuming Replication

## Preparing the RVG Volumes for Snapshot Operation

You must prepare the volumes under the RVG for the snapshot operation. For more information on preparing volumes, refer to the section "[Preparing the Volumes](#)" on page 159.

## Creating the Cache Object

If you intend to create instant space-optimized snapshots then you must create the cache object within the same disk group as the data volumes. Use the `vxassist make` command to create the cache volume. After creating the cache volume, create the cache object using the `vxmake cache` command. This command allows you to set the `autogrow` option for the cache object which allows the cache object to grow automatically, if the size of the cache object is not enough for the incoming writes.

For example, to create the cache volume of size 1GB with a name `cache-vol` and with a mirrored layout, type the following command on `seattle`:

```
# vxassist -g hrdg make cache-vol 1g layout=mirror init=active
```

Now, you can create a cache object named `cache-obj` for the cache volume by typing the following command on `seattle`:

```
# vxmake -g hrdg cache cache-obj cachevolname=cache-vol \
  autogrow=on regionsize=128
```

However, you can also create the cache object by specifying a value for the `cachesize` parameter in the `vxrvsg snapshot` command. This command creates one cache object for every space-optimized snapshot. To create one cache object for all the space-optimized snapshots, you must create the cache object using the `vxassist make` command.

If an RVG consists of only VxVM ISP volumes, then the `cachesize` parameter cannot be used. In this case, use the `vxvoladm` command to create an ISP cache object:

```
# vxvoladm make <cache_object> <cachesize> type=cachevolume \
  init=active
```

## Freezing or Pausing Replication

For more information on freezing or pausing the replication, refer to [“Freezing or Pausing Replication”](#) on page 160.



## Creating Instant Space-Optimized Snapshots

To create a space-optimized snapshot for each data volume in an RVG, use the following command:

```
# vxrvrg -g diskgroup [-P prefix] -S snapshot rvg_name \  
  [instantfull=volume_list [syncing=yes|no]] \  
  [exclude=volume_list] [plexbreakoff=volume_list] \  
  [plexprefix=plex_prefix] {cache=cachename|cachesize=size} \  
  [comment="<comment>"]
```

Use the `vxrvrg snapshot` command with its attributes to specify the type of snapshot that you want to create. By default, all the volumes in the RVG are considered for the space-optimized snapshots. To exclude any of the volumes from being considered for the space-optimized snapshots, use one of the following attributes. Depending on the attribute specified with the `vxrvrg snapshot` command, appropriate snapshots of the volumes are created.

The attribute `instantso` need not be specified when the `-S` option is specified.

The attribute `instantfull` specifies a comma-separated list of volumes that need to be included when creating an instant full snapshot of the volumes in an RVG.

The attribute `exclude` specifies a comma-separated list of volumes that do not need to be considered for any kind of snapshots.

The attribute `cache` specifies a name for the cache object. However, even if you do not specify a name, you can still specify a size for the cache. The `cachesize` attribute specifies a default size for the cache object with respect to the source volume. These operations together create one cache object per snapshot volume.

You can specify the volume list for the attributes `instantfull`, `exclude` or `plexbreakoff` when creating the instant space-optimized snapshots. This results in some volumes having an instant full snapshot, some with an instant space-optimized snapshot, some of them with instant plex-breakoff and some being excluded. Any volumes in the RVG that are not specified in the `volume_lists` of the attributes `exclude`, `plexbreakoff`, or `instantfull` will be snapped in the same way as the specified snapshot type, which is the instant space-optimized snapshot.

### Example: Creating Space-Optimized Snapshots

This example describes the steps to create an instant space-optimized snapshot for the specified RVG:

1. Prepare the required volumes if the volumes have not been prepared already.

```
# vxsnap -g hrdg prepare hr_dv01
# vxsnap -g hrdg prepare hr_dv02
```

Perform this operation for all the data volumes in the RVG for which you intend to create snapshots.

2. You can create the cache volume and the cache object if you want to create all the space-optimized snapshots on a single cache object. To do this, refer to [“Creating the Cache Object”](#) on page 165.

However, if you want to create separate cache objects for each snapshot proceed to the next step. You can create the cache object for each snapshot by specifying the *cachesize* or *cache* parameter.

3. Follow one of the steps provided depending on the method you have chosen for the cache object creation.

- ◆ To create the space-optimized snapshot for the volumes with a precreated cache object, issue the command:

```
# vxrvrg -g hrdg -S -P SO snapshot hr_rvg cache=snap-cacheobj
```

- ◆ To create the space-optimized snapshot for the volumes with a separate cache object for each volume, issue the command:

```
# vxrvrg -g hrdg -S -P S01 snapshot hr_rvg cachesize=10%
```

The cache objects are created for each snapshot with cache volumes that are 10% of the source volume. You can also specify an absolute value for the *cachesize* parameter.

---

**Note** If the size of the cache volume is less than 5MB, this command will fail.

---

### Unfreezing or Resuming Replication

After taking a snapshot, unfreeze replication if you are using IBC messaging; otherwise, if you have paused replication, resume it. The snapshots are now ready for use.



## Instant Plex-breakoff Snapshots

The `vxxrvg snapshot` command creates instant plex-breakoff snapshots of all the volumes in the RVG at a single point in time.

The steps required to create plex-breakoff snapshots are as follows:

- ◆ Preparing the RVG volumes for snapshot operation
- ◆ Creating snapshot plexes for Data Volumes in an RVG
- ◆ Freezing or Pausing Replication
- ◆ Creating an instant plex-breakoff snapshot
- ◆ Unfreezing or Resuming Replication

### Preparing the RVG Volumes for Snapshot Operation

It is necessary to prepare the volumes under the RVG for the snapshot operation. For more information on preparing volumes, refer to “[Preparing the Volumes](#)” on page 159.

### Creating Snapshot Plexes for Data Volumes in an RVG

You must create plexes for the required volumes before you take an instant plex breakoff snapshot.

Use the `vxsnap addmir` command to add one or more plexes to a volume:

```
# vxsnap -g diskgroup [-b] addmir volume [nmirror=<N>] \  
  [attributes...]
```

---

**Note** Run this command on every data volume in the RVG that needs the plex-breakoff snapshot to be created.

---

For more information on creating the plexes, refer to the *VERITAS Volume Manager Administrator's Guide*.

### Freezing or Pausing Replication

For more information on freezing or pausing the replication, refer to “[Freezing or Pausing Replication](#)” on page 160.



## Creating Instant Plex Breakoff Snapshots

The instant plex-breakoff snapshot feature enables you to create plex-breakoff snapshots just like the traditional snapshot feature. This operation cannot be used for an RVG that consists of VxVM ISP volumes.

### Prerequisites for creating instant plex breakoff snapshots

- ✓ Make sure the volumes for which you want to create plex-breakoff snapshots already have the appropriate plexes created and are in an SNAPDONE state.
- ✓ Make sure you create the plexes using appropriate prefixes in case you want to use specific plexes 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 of each data volume in an RVG, use the following command:

```
# vxrvrg -g diskgroup [-P prefix] snapshot rvg_name \
  [instantfull=volume_list [syncing=yes|no]] \
  [instantso=volume_list {cache=cachename|cachesize=size}] \
  [exclude=volume_list] [plexprefix=plex_prefix] \
  [comment="<comment>"]
```

Use the `vxrvrg snapshot` command with its attributes to specify the type of snapshot that you want to create. This is the default if neither the `-S` nor the `-F` option is specified. By default, all the volumes will be included for instant plex breakoff snapshots, provided that they have the plex volumes appropriately created. To exclude any of the volumes, use one of the following attributes. Depending on the attribute specified with the `vxrvrg snapshot` command, appropriate snapshots of the volumes are created.

The attribute `exclude` specifies a comma-separated list of volumes that do not need to be considered for any kind of snapshot.

The `plexprefix` attribute specifies a prefix for the plexes that will be used for creating the plex-breakoff snapshots. This is allowed only if the `-F` or `-S` option is not specified or if you have specified a list of volumes for creating plex-breakoff volumes with the `vxrvrg snapshot` command.



### Example:

This example describes the steps to create an instant plex breakoff snapshot for an RVG:

1. Prepare the required volumes if they have not been prepared already.

```
# vxsnap -g hrdg prepare hr_dv01
```

2. If the volumes for which the plex-breakoff snapshots need to be created do not have the required plexes, create them using the following command:

```
# vxsnap -g hrdg addmir hr_dv01
```

Repeat this step for all the required data volumes in the RVG. Initial synchronizing of the plexes may take some time depending on the size of the volume.

If you need to use specific plexes during the snapshot operation, make sure you name them appropriately when creating them. However, you can also do it later using the following command:

```
# vxedit -g hrdg rename hr_dv01-02 snapplex-dv01
```

3. Use the following command to create snapshots using the specific plex prefixes:

```
# vxrvg -g hrdg -P JULY snapshot hr_rvg plexprefix=snapplex
```

Use the following command to create the snapshots without specifying the plex prefix:

```
# vxrvg -g hrdg -P JULY1 snapshot hr_rvg
```

## Unfreezing or Resuming Replication

After taking a snapshot, unfreeze replication if you are using IBC messaging; otherwise, if you have paused replication, resume it. The snapshots are now ready for use.



## Administering Snapshots

### Refreshing Snapshots

The `vrxvg snaprefresh` command allows you to refresh the snapshots of the volumes in an RVG. It creates a new point-in-time image of the volume. For example, a snapshot taken on Monday can be refreshed on Tuesday. Before refreshing the snapshot make sure the data volumes are consistent at the application level by either freezing or pausing replication. After refreshing the snapshot, unfreeze replication if you are using IBC messaging; otherwise, if you have paused replication, resume it.

---

**Note** If the Secondary RVG is inconsistent, then VVR does not allow you to refresh the snapshots using the volumes under this RVG.

---

Use the following command to refresh existing snapshots:

```
# vrxvg -g diskgroup [-P <prefix>] snaprefresh rvg_name
```

---

**Note** After refreshing the snapshots you must issue the command `vxsnap syncstart` to start synchronizing the instant full snapshots. This is not required for instant space-optimized snapshots.

---

### Reattaching the Snapshot Plexes to the Data Volumes (snapback)

The `snapback` operation reattaches the snapshots of the instant full snapshot volume or the plexes of the plex breakoff snapshot volume back to the original volume. After working with the snapshot volumes, you can reattach the plexes to the data volumes in the RVG using the `snapback` operation. The `snapback` operation is instant as the plexes are resynchronized in the background.

---

**Note** The `snapback` operation can be performed only on instant full snapshots and plex-breakoff snapshots but not on space-optimized snapshots.

---

The `vrxvg snapback` command snaps back the snapshot, that is, reattaches the snapshot plexes to their respective data volumes in the RVG. Note that if an RVG consists of VxVM ISP volumes then you must use the `-f force` option to reattach the snapshot plexes to their data volumes. Not using the `force` option causes the operation to fail with an error message.

You can use the default action of the `vrxvg snapback` command if the data volumes in the RVG have only one snapshot. If the data volumes have more than one snapshot plex, use the `-a` option with the `vrxvg snapback` command to reattach all snapshots; the `-a` option snaps back all the plexes to their original data volumes.



Note that if the plexes had been added using the `vxsnap addmir` command, then `vxrvrg snapback` command will reattach the plexes in the `SNAPDONE` state. Otherwise, it will reattach the plexes in the `ACTIVE` state.

For example, use the `-a` option to reattach the snapshot volumes `june-dv_name` and `july-dv_name` to each data volume `dv_name` in the RVG.

The `-P` option when used with the `vxrvrg snapback` command enables you to reattach a specific set of snapshots that have been identified by the prefix. To snapback all the data volumes with the prefix `month` in their names, specify the prefix `month` using the `-P` option.

To snapback a single snapshot plex to the data volume in an RVG, use the following command:

```
# vxrvrg -g diskgroup -P prefix snapback rvg_name
```

To snapback the plexes from all snapshots, of the volume in an RVG, use the following command:

```
# vxrvrg -g diskgroup -a snapback rvg_name
```

All the snapshot plexes are snapped back to their original data volumes in the RVG `rvg_name`.

To snapback the snapshot plexes with a specific prefix for all data volumes in an RVG, use the following command:

```
# vxrvrg -g diskgroup -P june snapback rvg_name
```

Snapshot plexes with the name `june-dv_name` are snapped back to their original data volumes in the RVG.

## Restoring Data from the Snapshots

Use the `vxxvrg snaprestore` command to restore the data from the required snapshots.

---

**Note** When restoring the volumes you must specify the exact snapshot corresponding to the respective volumes that need to be restored. Therefore, it is recommended that you create snapshots with prefixes using the `-P` option so that they can be easily restored. However, if you do not specify any prefix, the default prefix `SNAP` will be used. The `vxxvrg snaprestore` command can be used only for restoring the data from the instant snapshots.

---

```
# vxxvrg -g diskgroup [-P prefix] snaprestore rvg_name
```

The restore operation is very useful when the data in a volume is corrupted, for example, because of a software or user error, and must be restored to a known state preserved in a snapshot of that volume taken some time earlier. Because a replicated volume under the RVG propagates all changes to the Secondary, the Secondary volume must be restored back to a known state. VVR can now do this automatically using the instant snapshots. It uses the DCM logging to resynchronize the Secondary so that only the regions changed as part of the restore operation are applied to the Secondary volumes.

If there are multiple RLINKs in an RVG, then VVR synchronizes all the Secondary hosts using the bits on the DCM log. If one of the RLINKs is already in a DCM mode, then VVR also updates the bits corresponding to the regions that need to be restored as a part of the restore operation to the DCM. Now, VVR resynchronizes all the Secondary hosts using the consolidated bits on the DCM log.

---

**Note** In case of the multiple RLINK setup, if either the autosync or resync operation was already in progress across some RLINK in the RVG, then the resynchronization for the other RLINKs that have already switched to the DCM mode as a part of the restore operation starts automatically.

---

The `vxxvrg snaprestore` command can be used to restore data both from the Primary and the Secondary. On the Primary the `vxxvrg snaprestore` command populates the DCM for replay that can be used on the Secondary only if it has no attached RLINKS or if the RLINK is in the FAIL state. For more information, refer to [“Rules for restoring volumes.”](#)

---

**Note** Restoring data on the Primary RVG volumes using the `vxxvrg snaprestore` command deletes all the existing checkpoints.

---



## Rules for restoring volumes

The volumes in an RVG can be restored only according to the rules mentioned below. This is irrespective of whether the volumes are restored using the `vrxvg snaprestore` command or specific volumes under the RVG are restored using the `vxsnap restore` command.

### On Primary

1. If the RLINK is detached, the volume is restored like any other VxVM volume.
2. If the RLINK is active, the RLINK is put into DCM logging mode and the regions that need to be modified by the restore operation are marked on the DCM, and the volume is restored. The RVG must be resynchronized using `vrxvg resync` command, to ensure that the restored data is available on the Secondary RVG. This is independent of the SRL protection setting and works even if the `srlprot` attribute is not set to `dcm` or `autodcm`.
3. If the RLINK is active and the volumes are not DCM logging enabled, then the restore operation fails, unless the `-f` (force) option is specified. If the force option is specified, the RLINK is detached before the volume is restored.

### On Secondary

The restore operation is allowed only if:

- ◆ the RLINKs are detached.
- ◆ the attached RLINK in the RVG is in the FAIL state.

## ▼ To restore the snapshots on the Primary

### On Primary:

1. To stop the specific RVG use the following command:  

```
# vrxvg -g hrdg stop hr_rvg
```
2. To restore the volumes from snapshot with a specific prefix, use the following command:  

```
# vrxvg -g hrdg -P JULY snaprestore hr_rvg
```

The RLINK changes to the DCM mode if it is not already in this mode.
3. To replicate the new changes to the Secondary, use the following command:  

```
# vrxvg -g hrdg resync hr_rvg
```

## Displaying the Snapshot Information

The `vxrvrg snapprint` command displays information on the relationship that exists between the original volumes and the corresponding snapshots. To display information on the snapshots use the following command:

```
# vxrvrg -g diskgroup snapprint rvg_name
```

The output of the command resembles:

```
vxrvrg snapprint hr_rvg
Creation Time : Fri Feb 14 02:25:58 2003
Source Volume      Snapshot Volume      Snapshot Type      Sync Status
-----
hr-dv01            JULY1-dv01           Inst-Full          Complete
hr-dv02            JULY1-dv02           Inst-Full          Complete

Creation Time : Fri Feb 14 02:25:45 2003
Source Volume      Snapshot Volume      Snapshot Type      Sync Status
-----
hr-dv01            JULY-dv01            Inst-Full          Complete
hr-dv02            JULY-dv02            Inst-Full          Complete

Creation Time : Fri Feb 14 01:46:38 2003
Source Volume      Snapshot Volume      Snapshot Type      Sync Status
-----
hr-dv01            S01-dv01             Inst-SO            Incomplete
hr-dv02            S01-dv02             Inst-SO            Incomplete

Creation Time : Fri Feb 14 01:44:55 2003
Source Volume      Snapshot Volume      Snapshot Type      Sync Status
-----
hr-dv01            SO-dv01              Inst-SO            Incomplete
hr-dv02            SO-dv02              Inst-SO            Incomplete

Creation Time : Thu Feb 13 09:14:11 2003
Source Volume      Snapshot Volume      Snapshot Type      Sync Status
-----
hr-dv01            JUNE-vol1            Inst-Full          Complete
hr-dv02            JUNE-vol2            Inst-Full          Complete
```

---

**Note** The `vxrvrg snapprint` command can also be used to display the status of the snapshots that have been created using the traditional snapshot feature. However, this output will not display the correct time.

---



## Destroying the Snapshots

The `vxrvg snapdestroy` command enables you to destroy or delete the snapshot volumes from the RVG. The `vxrvg snapdestroy` command first dissociates the snapshot volumes from the original volumes and then destroys the volumes.

To destroy the snapshot volumes, use the following command:

```
# vxrvg -g diskgroup [-P prefix] [-o keepcache] snapdestroy rvg_name
```

The argument `snapdestroy` along with the different attributes specifies the snapshot that is to be destroyed.

By default, the `snapdestroy` attribute removes the cache object along with the instant snapshots of the specified prefix. However, if you specify the `-o keepcache` option, then the cache object is not deleted. The `-o keepcache` option can be specified only for the pre-created cache objects. The same cache object can then be used for creating new snapshots.



## Using the Traditional Snapshot Feature

This snapshot feature of VVR enables you to break off mirrors from the data volumes in an RVG thus providing snapshots of the data volumes in the RVG. Snapshots can be used to perform operations such as Decision Support Systems (DSS) and backup. Snapshots can also be used to retain a consistent copy of the Secondary data volumes during Data Change Map (DCM) resynchronization.

The `vxrv snapshot` command takes a snapshot of all the volumes in the RVG at a single point in time; therefore, the operation is atomic in nature. The `vxrv snapback` command reattaches the plexes of the snapshot volumes to the original data volumes in the RVG. The `vxrv snapshot` command creates the same type of snapshot on the data volumes as a `vxassist snapshot` command would create on a volume associated or unassociated with an RVG. To snapshot and snapback a specific volume or specific plexes of one or more volumes, use the `vxassist` command.

If an RVG contains component volumes of a volume set, the `vxrv snapshot` command cannot be used to take snapshots of its data volumes.

Using the snapshot feature involves the following tasks:

- ◆ [Creating Snapshot Plexes for Data Volumes in an RVG](#)
- ◆ [Freezing or Pausing Replication](#)
- ◆ [Taking a Snapshot](#)
- ◆ [Unfreezing or Resuming Replication](#)
- ◆ [Reattaching the Snapshot Plexes to the Data Volumes \(Snapback\)](#)

### Creating Snapshot Plexes for Data Volumes in an RVG

To use the RVG snapshot feature, create snapshot plexes for each data volume in the RVG. Creating the snapshot plexes is a one-time operation.

To create a snapshot plex for a volume, use the following command:

```
# vxassist -g diskgroup snapstart dv_name
```

The `vxassist snapstart` command creates a new plex for the volume `dv_name` and attaches it to the volume. When the attach is complete, the state of the plex is `SNAPDONE` and a snapshot can be taken.



## Freezing or Pausing Replication

Before taking the snapshot on the Secondary, make the data volumes consistent at the application level by either freezing or pausing replication. To make the data volumes consistent at the application level, use the IBC Messaging utility `vxibc`. For more information, see “[Using VVR for Off-Host Processing](#)” on page 199. For a failed Primary, pause the Primary RLINK, and then take a snapshot of the RVG. If you do not use `vxibc`, pause the RLINK before taking the snapshot.

## Taking a Snapshot

The `vrxvg snapshot` command takes snapshots of the data volumes in the RVG. It creates a snapshot volume with the name `SNAP-dv_name` for each data volume in the RVG.

Each data volume in an RVG can have more than one snapshot volume. The `-P` option to the `vrxvg snapshot` command 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`.

To take a snapshot of each data volume in an RVG, use the following command:

```
# vxrvg -g diskgroup snapshot rvg_name
```

To specify a prefix for the snapshot of each data volume in an RVG, use the following command:

```
# vxrvg -g diskgroup -P june snapshot rvg_name
```

A snapshot data volume with the name `june-dv_name` is created for each data volume in the RVG. You can have more than one snapshot of the data volumes in an RVG.

Perform the required operation on the snapshots; then `snapback`, that is, reattach the snapshots to the data volumes in the RVG using the `vrxvg snapback` command.

## Unfreezing or Resuming Replication

After taking a snapshot, unfreeze replication if you are using IBC messaging; otherwise if you have paused replication resume it. The snapshots are ready for use.

## Performing the Required Operations on the Snapshot

Use snapshots to perform off-host processing operations including Decision Support Systems (DSS), backup, and trial failover in VVR. Snapshots can also be used to keep a consistent copy of the data volumes in an RVG when DCM resynchronization is in progress. After performing the required operation on the snapshots, reattach them.

## Reattaching the Snapshot Plexes to the Data Volumes (Snapback)

The snapback operation reattaches a snapshot volume with the original volume. After working with the snapshot volumes, reattach them to the data volumes in the RVG. The snapback operation may take a long time to complete because it performs a complete resynchronization of the snapshot plexes. To perform a faster and more efficient snapback operation, see [“Using the Traditional Snapshot Feature”](#) on page 177.

The `vrxvg snapback` command snaps back, that is, it reattaches the snapshot plexes to their respective data volumes in the RVG.

You can use the default action of the `vrxvg snapback` command if the data volumes in the RVG have one snapshot. If the data volumes have more than one snapshot plex, use the `-a` option with the `vrxvg snapback` command to reattach all snapshots; the `-a` option snaps back all the plexes to their original data volumes. For example, use the `-a` option to reattach the snapshot volumes `june-dv_name` and `july-dv_name` to each data volume `dv_name` in the RVG.

The `-P` option to the `vrxvg snapback` command enables you to reattach a specified snapshot. To reattach all the data volumes with the prefix `month` in their names, specify the prefix `month` using the `-P` option.

For data volumes with single snapshot plexes in an RVG, snapback using the following command:

```
# vrxvg -g diskgroup snapback rvg_name
```

To snapback all plexes for each data volume in an RVG, use the following command:

```
# vrxvg -g diskgroup -a snapback rvg_name
```

All the snapshot plexes are snapped back to their original data volumes in the RVG `rvg_name`.

To snapback snapshot plexes with a specific prefix for all data volumes in an RVG, use the following command:

```
# vrxvg -g diskgroup -P june snapback rvg_name
```

Snapshot plexes with the name `june-dv_name` are snapped back to their original data volumes in the RVG.



To snapback a plex to a specific data volume in an RVG, use the following command:

```
# vxassist -g diskgroup snapback SNAP-dv_name
```

For more information on using the `vxassist snapback` command, see the *VERITAS Volume Manager Administrator's Guide*.

## Using Snapback with `resyncfromreplica` Option

The default action of the `vxassist snapback` command is to resynchronize the snapshot plex with the contents of the original volume. The `resyncfromreplica` option of the `vxassist snapback` command synchronizes the original volume with the contents of the snapshot plex. This operation is similar to a restore from a backup operation. In most cases, the default action of the `vxassist snapback` command must be used, but there are some situations where the `resyncfromreplica` option may be used. The `vxrvg snapback` command does not provide the `resyncfromreplica` option, therefore, the operation must be performed one volume at a time. The `resyncfromreplica` operation is not allowed on the Primary or the Secondary SRL.

---

**Caution** Improper use of the `resyncfromreplica` option of the `vxassist snapback` command on a replicated volume can cause data corruption. You must read the following sections before proceeding.

---

### Using `resyncfromreplica` Option to Recover from Logical Corruption of Data

If there is a logical corruption of data and a good snapshot of the data volumes exists, it can be used to restore the data volumes to a version before the error occurred. If the snapshot exists on the Primary, before issuing the `vxassist -o resyncfromreplica snapback` command, shutdown the application and detach all the RLINKs. The `resyncfromreplica` operation will fail if the RLINK is not detached. On completing the snapback operation, perform a complete synchronization of the Secondary data volumes. Refer to [“Methods to Synchronize the Secondary”](#) on page 58.

If the snapshots exist on the Secondary, before issuing the `vxassist -o resyncfromreplica snapback` command, migrate the Primary role to this Secondary host, but do not start the application. Refer to [“Migrating the Primary”](#) on page 216. After migrating the Primary role, detach the RLINK to the original Primary, which is now a Secondary, and then perform the snapback operation. On completing the snapback operation, perform a complete synchronization of the Secondary data volumes.

If you choose to completely synchronize the Secondary using a checkpoint, make sure that any Primary checkpoints that were taken before the snapback operation are not used to resynchronize the Secondary. VVR may show these checkpoints as valid if they have not

overflowed; however, the checkpoints are not valid. You can only use Primary checkpoints taken after the `resyncfromreplica` operation to resynchronize the Secondaries.

### Using `resyncfromreplica` to Recover Failed Secondary Data Volumes

The `resyncfromreplica` option can also be used to restore Secondary data volumes that are corrupt due to disk errors. In this case, the data volumes can be restored from existing snapshots. The RLINK must be in the FAIL state to perform the `resyncfromreplica` operation. See “[Restoring the Secondary from Online Backup](#)” on page 192. Use these snapshots instead of backups.

If you choose to restore the Secondary using checkpoints, you must ensure that:

- ◆ The snapshot volume being used for the `resyncfromreplica` operation corresponds to the checkpoint to be used in the `vxrlink restore` command for the RLINK.
- ◆ The checkpoint is still valid before proceeding with the `resyncfromreplica` snapback operation. Issue the following command to determine whether the checkpoint is still valid:

```
# vxrlink -g diskgroup cplist rlink_name
```

- ◆ The snapshot volumes were never written to.

VVR cannot ensure or check if the above conditions are met and failure to meet all of the conditions can result in an inconsistency between the Primary and the Secondary.



## Using VERITAS Volume Manager FastResync

FastResync (FR) enables you to split off a plex from a mirrored volume, manipulate it, and then reattach it to the original volume without doing a complete resynchronization of the volume. FastResync is a separately licensed feature of VxVM.

FR maintains a bitmap of changes to the volume while the plex was split off, and also of changes to the split-off mirror. When the plex is attached, only the blocks represented in the map are resynchronized.

You can use FR to perform the snapback operation after the off-host processing operation is complete.

The operations performed by the `vxrvg snapshot` and `vxrvg snapback` commands are based on one of the following conditions:

- ◆ If you do not have an FR license, the `vxrvg snapshot` command creates a simple snapshot without any FR bitmaps. A `vxrvg snapback` operation then results in a full resynchronization of the plexes.
- ◆ If you have an FR license and FastResync is enabled on the volumes, but no DCO logs are attached to the data volumes, the `vxrvg snapshot` command creates a snapshot with a non-persistent FR bitmap. A `vxrvg snapback` operation performs a FastResync, but if the system reboots at any time after the snapshot operation, the information in the FR bitmap is lost and a full resynchronization takes place.
- ◆ If you have an FR license, FastResync is enabled, and DCO logs are attached to the data volumes, the `vxrvg snapshot` command creates a snapshot with persistent FR bitmap. A `vxrvg snapback` operation performs a FastResync even if the system reboots at any point in time after the snapshot operation.

For more information on Persistent and Non-Persistent FR, see the *VERITAS Volume Manager Administrator's Guide*.

### Enabling FastResync

To enable FR on a data volume, type:

```
# vxvol -g diskgroup set fmr=on dv_name
```

Do this for all the data volumes in the RVG that you want to access.

Refer to the *VERITAS Volume Manager Administrator's Guide* for more information.

## Verifying the DR Readiness of a VVR Setup

When setting up a Disaster Recovery (DR) solution, it is very important to verify the effectiveness of the DR solution. Although VVR guarantees that integrity of data is maintained between the Primary and Secondary data volumes, validating data is necessary to ensure that there has been no data loss due to administrative error, user error, or some other technical reasons. Validation also helps you to be certain that the data that has been replicated to the Secondary (Disaster Recovery site) can be used to bring up the applications in case of a disaster.

The way to validate the DR readiness of the DR site is to bring up the application on the DR site. This can be done in two ways. One is to migrate the Primary role to the Secondary and then run the applications on the new Secondary using the replicated data. Another way of performing a firedrill is using the snapshot feature. Using this feature VVR creates snapshots of the data volumes that can be used to bring up the application on the Secondary.

Data validation can be used to verify the integrity of the data that has been replicated to the Secondary from the Primary. This is done by comparing the data with the data on the Primary. When the Secondary data volumes are validated after the replication has been stopped, the volumes are dissociated from an RVG. This could be very useful in case you want to validate the data volume before adding it back to the RDS. However, data can also be validated online, that is, when the replication is in progress. This is achieved by using the instant space-optimized snapshot feature to create point in time snapshots of the primary and secondary data volumes. In this case, instead of the actual volumes the snapshot volumes are compared and validated. For information on the methods to create the snapshots, refer to [“Creating RVG Snapshots”](#) on page 157.

VVR enables you to validate the DR Readiness of the Secondary by using one of the following methods.

- ◆ [Performing a Failover](#)
- ◆ [Performing a Firedrill](#)
- ◆ [Verifying the Data on the Secondary](#)

### Performing a Failover

A disaster like scenario can be tested by using the migrate operation to perform a complete failover testing. This can be done by migrating the role of the Secondary to a Primary and making sure that the application is running on the new Primary. For more information on how we can perform the role transfer refer to the chapter [“Transferring the Primary Role”](#) on page 215.



## Performing a Firedrill

Firedrill is a process of bringing up the application on the Secondary using the replicated data. This data is then used to perform some processing in order to validate the consistency and correctness of the data.

To test the failover you can use a point-in-time image of the data on the Secondary data volumes. VVR provides you with the option of creating instant full and space-optimized snapshots. For any information on the methods to create snapshots see section “[Creating RVG Snapshots](#)” on page 157. You can use the appropriate type of snapshot method to create the snapshots. The instant space-optimized snapshot has much lesser space requirement compared to the instant full or plex-breakoff snapshots. These space-optimized snapshots are used to test the Secondary failover.

### Points to Note

- ◆ A firedrill cannot be performed using the Secondary volumes therefore the snapshots must be used.
- ◆ The secondary must be in a consistent state when the snapshots are taken.
- ◆ When performing a firedrill no IBC messages are required to be sent so that a failover scenario similar to a real one is simulated.

### Automating the Firedrill Procedure

The firedrill procedure is most effective only when it is performed on a regular basis. The above method requires you to test the Secondary failover, manually, at frequent intervals. However, in case VVR is used in a VCS setup where the appropriate agents are installed, then the firedrill procedure can be automated using the `RVGSnapshot` and `RVGPrimary` agents that VCS provides. For more information on how you can use these agents to automate the firedrill testing, refer to the VCS Documents.



## Verifying the Data on the Secondary

VVR enables you to verify that the data on the Secondary is identical to the data on the Primary data volumes, either when the application is active or inactive. VVR provides the following methods to verify the data at the Secondary site: online data verification and offline data verification.

Online data verification allows you to validate the data even when 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 replication is not active. If you have already created the Primary and Secondary volumes and replication is in progress, you need to pause replication and then perform data validation between the corresponding Primary and Secondary volumes to make sure that the data volumes on the Primary and Secondary are the same. To do this, use the `vradmin syncrvg` command with the `-verify` option. To use this command for verifying the data, the Secondary must be up-to-date. This command performs a comparison of the checksums on the corresponding Primary and Secondary volumes.

You can also validate the data on new data volumes before adding them to an RDS. For information, refer to [“Verifying the Data on the Primary and Secondary Volumes”](#) on page 123.

### Performing Online Data Verification

The space-optimized snapshots that are created using the `vxrvg snapshot` command 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 `vradmin -verify syncrvg` command 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 very useful if you want to check the integrity of the data volumes on the Secondary when replication is in progress.

The `vradmin verifydata` command creates the space-optimized snapshots on the Primary and the Secondary before it proceeds with performing online data verification. The `vradmin verifydata` command 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 `vradmin verifydata` then verifies the data between the remote and local hosts by comparing the space-optimized snapshots.



To perform online data verification, use the command:

```
vradmin [-g diskgroup] [-k {cache|snap}] verifydata rvg_name \  
      sechost {cache=cacheobj/cachesize=size}
```

The attribute *sechost* specifies the name of the Secondary host.

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. If the RVG on either the Primary or the Secondary has VxVM ISP volumes, then you cannot use the *cachesize* attribute.

You must specify only one of these attributes at one time for the command to create a cache object for each snapshot.

This command performs the following tasks:

- Registering the application on the Primary and the Secondary.
- Freezing replication on the Primary and Secondary.
- Taking snapshots and verifying the data.
- Destroying the snapshots.

By default, the `vradmin verifydata` command destroys the snapshot volume and the cache object after the data verification has proceeded successfully. However, if you want to preserve the snapshot volumes then you must use the `vradmin verifydata` command with the `-k snap` option. If you want to preserve the cache object then use the `vradmin verifydata` command with the `-k cache` option. The same cache object can then be reused when creating future snapshots. You cannot use the `-k` option if you have used the `cachesize` option, as it is an invalid combination and the command fails with an error message. Note that when specifying the `-k` option you must specify either the *cache* or the *snap* argument with it.

---

**Note** When the `-k snap` option is specified the cache object is also preserved along with the snapshot since the snapshot cannot exist without the cache object.

---

VVR also provides you with sample scripts that can be used to freeze the replication and then take instant space-optimized snapshots. For details, refer to [“Sample Scripts”](#) on page 210.

1. Prepare the volumes that need to be included in the snapshot. For more information on preparing volumes, refer to [“Preparing the Volumes”](#) on page 159.
2. Create the required cache object within the same disk group as the data volume. For more details on creating the cache object, refer to [“Preparing the RVG Volumes for Snapshot Operation”](#) on page 164.

- Use the following command to perform online data verification:

```
vradmin [-g diskgroup] [-k {cache|snap}] verifydata rvg_name \  
      sechost {cache=cacheobj | cache=size}
```

## Performing Offline Data Verification

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 `vradmin syncrvg` command with the `-verify` option verifies and reports any differences between the data volumes associated with the Secondary RVG and the corresponding Primary RVG. The `vradmin -verify syncrvg` command only reports whether the Primary and Secondary volumes are identical or not. It does not make them identical. As the command runs, it reports the progress every 10 seconds. An MD5 checksum is used to calculate the difference between the Primary and the Secondary data volumes. Refer to [“Using Difference-Based Synchronization”](#) on page 70.

Prerequisite for using the `vradmin -verify syncrvg` command:

- ✓ All applications using the Primary data volumes must be stopped before running the `vradmin -verify syncrvg` command.

### ▼ To verify the differences between the Primary and Secondary data volumes

```
# vradmin -g diskgroup -verify syncrvg local_rvgname sec_hostname...
```

When this command is invoked, you are prompted to confirm that the Primary data volumes are not in use. You can use the `-s` option to skip this confirmation step.

The argument `local_rvgname` is the name of the RVG on the local host and represents the RDS.

The argument `sec_hostname` is a space-separated list of the names of the Secondary hosts as displayed in the output of the `vradmin printrvg` command.

This command checks the status of the Primary RLINK to each of the Secondary RVGs being verified. If any of the RLINKs are not up-to-date, the `vradmin -verify syncrvg` command returns with a message to indicate that the RLINKs are not up-to-date. In this scenario, verification is not performed. Use the `vxlink status` command to determine the extent to which the Secondary is behind.



**Example:**

To verify the data differences between the Primary RVG `hr_rvg` on `seattle` and the Secondary RVG on host `london`, issue the following command from any host in the RDS:

```
# vradmin -g hrdg -verify syncrvg hr_rvg london
```

The output resembles the following if the Primary and Secondary data volumes are identical:

```

Message from Primary:
VxVM VVR vxrsync INFO V-5-52-2210 Starting volume verification to remote
VxVM VVR vxrsync INFO V-5-52-2211 Source host: 10.182.136.192
VxVM VVR vxrsync INFO V-5-52-2212 Destination host(s): 10.182.136.193
VxVM VVR vxrsync INFO V-5-52-2213 Total volumes: 1
VxVM VVR vxrsync INFO V-5-52-2214 Total size: 4.000 G

Eps_time Dest_host Src_vol Dest_vol F'shed/Tot_sz Diff Done
00:00:00 10.182.136.193 hr_dv hr_dv 0M/4096M 0% 0%
00:00:10 10.182.136.193 hr_dv hr_dv 221M/4096M 0% 5%
Message from Primary:
00:00:20 10.182.136.193 hr_dv hr_dv 468M/4096M 0% 11%
Message from Primary:
00:00:30 10.182.136.193 hr_dv hr_dv 705M/4096M 0% 17%
Message from Primary:
00:00:40 10.182.136.193 hr_dv hr_dv 945M/4096M 0% 23%
Message from Primary:
00:00:50 10.182.136.193 hr_dv hr_dv 1184M/4096M 0% 29%
Message from Primary:
00:01:00 10.182.136.193 hr_dv hr_dv 1419M/4096M 0% 35%
Message from Primary:
00:01:10 10.182.136.193 hr_dv hr_dv 1655M/4096M 0% 40%
Message from Primary:
00:01:20 10.182.136.193 hr_dv hr_dv 1886M/4096M 0% 46%
Message from Primary:
00:01:30 10.182.136.193 hr_dv hr_dv 2124M/4096M 0% 52%
Message from Primary:
00:01:40 10.182.136.193 hr_dv hr_dv 2356M/4096M 0% 58%
00:01:50 10.182.136.193 hr_dv hr_dv 2590M/4096M 0% 63%
Message from Primary:
00:02:00 10.182.136.193 hr_dv hr_dv 2838M/4096M 0% 69%
Message from Primary:
00:02:10 10.182.136.193 hr_dv hr_dv 3091M/4096M 0% 75%
Message from Primary:
00:02:20 10.182.136.193 hr_dv hr_dv 3324M/4096M 0% 81%
Message from Primary:
00:02:30 10.182.136.193 hr_dv hr_dv 3564M/4096M 0% 87%
Message from Primary:
00:02:40 10.182.136.193 hr_dv hr_dv 3809M/4096M 0% 93%
Message from Primary:
00:02:50 10.182.136.193 hr_dv hr_dv 4070M/4096M 0% 99%
00:02:51 10.182.136.193 hr_dv hr_dv 4096M/4096M 0% 100%
VxVM VVR vxrsync INFO V-5-52-2217 The volumes are verified as identical.

VxVM VVR vxrsync INFO V-5-52-2219 VxRSync operation completed.
VxVM VVR vxrsync INFO V-5-52-2220 Total elapsed time: 0:02:51

```



If there are differences in the data volumes, the output looks similar to the one shown below:

```

Message from Primary:
VxVM VVR vxrsync INFO V-5-52-2210 Starting volume verification to remote
VxVM VVR vxrsync INFO V-5-52-2211 Source host: 10.182.136.192
VxVM VVR vxrsync INFO V-5-52-2212 Destination host(s): 10.182.136.193
VxVM VVR vxrsync INFO V-5-52-2213 Total volumes: 1
VxVM VVR vxrsync INFO V-5-52-2214 Total size: 4.000 G

Eps_time Dest_host Src_vol Dest_vol F'shed/Tot_sz Diff Done
00:00:01 10.182.136.193 hr_dv hr_dv 0M/4096M 0% 0%
00:00:11 10.182.136.193 hr_dv hr_dv 231M/4096M 48% 6%
Message from Primary:
00:00:21 10.182.136.193 hr_dv hr_dv 476M/4096M 23% 12%
Message from Primary:
00:00:31 10.182.136.193 hr_dv hr_dv 719M/4096M 15% 18%
Message from Primary:
00:00:41 10.182.136.193 hr_dv hr_dv 954M/4096M 12% 23%
Message from Primary:
00:00:51 10.182.136.193 hr_dv hr_dv 1202M/4096M 9% 29%
Message from Primary:
00:01:01 10.182.136.193 hr_dv hr_dv 1438M/4096M 8% 35%
Message from Primary:
00:01:11 10.182.136.193 hr_dv hr_dv 1680M/4096M 7% 41%
Message from Primary:
00:01:21 10.182.136.193 hr_dv hr_dv 1924M/4096M 6% 47%
Message from Primary:
00:01:31 10.182.136.193 hr_dv hr_dv 2165M/4096M 5% 53%
Message from Primary:
00:01:41 10.182.136.193 hr_dv hr_dv 2418M/4096M 5% 59%
Message from Primary:
00:01:51 10.182.136.193 hr_dv hr_dv 2668M/4096M 4% 65%
00:02:01 10.182.136.193 hr_dv hr_dv 2906M/4096M 4% 71%
Message from Primary:
00:02:11 10.182.136.193 hr_dv hr_dv 3140M/4096M 4% 77%
Message from Primary:
00:02:21 10.182.136.193 hr_dv hr_dv 3386M/4096M 3% 83%
Message from Primary:
00:02:31 10.182.136.193 hr_dv hr_dv 3630M/4096M 3% 89%
Message from Primary:
00:02:41 10.182.136.193 hr_dv hr_dv 3881M/4096M 3% 95%
Message from Primary:
00:02:49 10.182.136.193 hr_dv hr_dv 4096M/4096M 3% 100%
VxVM VVR vxrsync INFO V-5-52-2218 Verification of the remote volumes found
differences.

VxVM VVR vxrsync INFO V-5-52-2219 VxRSync operation completed.
VxVM VVR vxrsync INFO V-5-52-2220 Total elapsed time: 0:02:50
    
```



## Backing Up the Secondary

You must take backups on the Secondary on a regular basis to guard against the consequences of disk failures. The Secondary checkpointing facility allows volume-level backups of the RVG to be restored at the Secondary node.

To get a consistent backup, replication must not be active. You do this by initiating a Secondary checkpoint at the Secondary node. This causes a request to be sent to the Primary node to pause updates and record the Secondary checkpoint in the SRL. While replication is paused, take a block-level backup of the RVG at the Secondary node. When the backup is complete, resume replication. Initiating a `resume` at the Secondary node causes a request to be sent to the Primary node to resume updates. Note that the Secondary cannot do a checkpoint if it has lost contact with the Primary.

If you must recover from Secondary data volume failure, then after the block-level backup has been restored, subsequent updates can be replayed from the Primary starting at the checkpoint and the Secondary can be brought up-to-date. The Secondary can be brought up-to-date only if the updates are still in the SRL. You can display a list of checkpoints on the Secondary using the `vxrlink cplist` command. For more information, see [“Displaying a List of Checkpoints”](#) on page 91.

## Checkpoint Pause/Resume Secondary RLINK

If a Secondary data volume fails and there is a checkpoint backup created as outlined above, you can restore from this backup copy without having to do a full Primary resynchronization of all the volumes. This procedure is also referred to as doing an online restore for the Secondary, because it is not necessary to stop the Primary RVG to bring a new copy of the Secondary data volume up-to-date.

---

**Note** The names of existing Secondary checkpoints can be obtained by performing the `vxrlink cplist` command on the Primary. The `vxrlink cplist` command can also be used to monitor whether the earlier checkpoints are about to overflow.

---

The checkpoint string can be up to 19 characters long for either a Primary or a Secondary checkpoint. Only the most recent Primary checkpoint string is displayed in the Primary RVG.

Unlike a simple Secondary pause, a checkpoint Secondary pause will fail if the Secondary is disconnected from the Primary at the time the command is executed, because communication with the Primary is required to create the checkpoint.



**▼ To create a Secondary checkpoint**

On the Secondary:

1. Pause the RLINK using a checkpoint.

```
# vxrlink -g diskgroup -c sec_checkpointname pause rlink_name
```

---

**Note** Pausing the RLINK with a checkpoint on the Secondary creates the Secondary checkpoint.

---

2. Back up all the volumes in the Secondary RVG, using a block-level backup.
3. Resume the RLINK.

```
# vxrlink -g diskgroup resume rlink_name
```

**▼ To delete a Secondary checkpoint**

1. Pause the RLINK using a checkpoint on Secondary.

```
# vxrlink -g diskgroup -c sec_checkpointname pause rlink_name
```

2. Delete the Secondary checkpoint, using the following command:

```
# vxrlink -g diskgroup -c sec_checkpointname checkdelete \  
rlink_name
```

---

**Note** Perform [step 2](#) only on the Primary.

---

## Restoring the Secondary from Online Backup

### Restoring from Secondary Checkpoints

If a Secondary volume becomes corrupted due to I/O errors, the volume can be restored from backup. When a `vxrlink restore` is initiated, a request is sent to the Primary to begin updates from a previously recorded checkpoint. A `restore` is not guaranteed to succeed though because checkpoints can become stale which means that the Primary has stopped maintaining the updates necessary for the restore. If this occurs, the Secondary RVG must be re-initialized using a Primary checkpoint or an `autosync attach` instead of being restored from backup.



## Restoring a Secondary RLINK

If a Secondary data volume fails the RLINK is put into the FAIL state. A restore from an online backup copy becomes necessary. This can only be done if a suitable Primary or Secondary checkpoint exists. If a Primary checkpoint still exists, it can be used if there is no Secondary checkpoint.

To restore a Secondary from an on-line backup, first restore the data from the on-line backup to all of the volumes. Because of internal constraints, you must restore all volumes even if only one has failed. (The normally read-only Secondary data volumes are writable while the Secondary is in FAIL state.) Then, execute the

`vxrlink -c checkpoint_name restore rlink` command, which causes the Secondary to request all updates which were made subsequent to the checkpoint from the Primary.

As with Primary checkpoints, if the checkpoint is not used before the SRL wraps around and the SRL overflows, the checkpoint will become STALE. If the checkpoint becomes STALE, you cannot use the methods described in this section to restore the data. You must synchronize the RLINK. See “[Methods to Synchronize the Secondary](#)” on page 58 for details. To prevent the checkpoint from becoming STALE, make sure the SRL is large enough to hold all of the updates which occur between the `vxrlink -c checkpoint pause` command and the `vxrlink -c checkpoint restore` command.

On the Secondary:

1. Assuming the RLINK is in the FAIL state, restore the backup to the data volumes.
2. Restore the RLINK to initiate updates to the Secondary data volumes:

```
# vxrlink -g diskgroup -c checkpoint_name restore rlink_name
```

---

**Note** In situations where you need to perform a restore when the RLINK is not in a FAIL state, use the following command to get the RLINK in a FAIL state:

```
# vxrlink -g diskgroup -w pause rlink_name
```

For example, you need to get an RLINK back in the FAIL state if a data volume fails, you restore from backup, and then after performing the restore command, you realize that the wrong backup was used. In this case, you need to get the RLINK back in the FAIL state before you perform the `restore` command.

---

While the restore is occurring, the RLINK is inconsistent. It becomes consistent when the `vxrlink restore` command completes successfully.



## Changing the VVR Tunables

VVR provides you with a number of tunable parameters that can be tuned to a specific value, according to your requirements. For a detailed explanation of the VVR tunables, see the *VERITAS Volume Replicator Planning and Tuning Guide*.

The following table provides you with a quick reference to the tunables that can be modified using the `vxtune` utility. All tunables can be modified using the system-specific method.

Tunable Name	Modifying Tunables Using the <code>vxtune</code> Utility	Auto-tuning Supported
<code>vol_rvio_maxpool_sz</code>	Yes	No
<code>vol_min_lowmem_sz</code>	Yes	Yes
<code>vol_max_rdback_sz</code>	Yes	No
<code>vol_max_nmpool_sz</code>	Yes	No
<code>vol_dcm_replay_sz</code>	No	No
<code>vol_nm_hb_timeout</code>	No	No
<code>voliomem_chunk_size</code>	No	No
<code>vol_vvr_use_nat</code>	No	No
<code>volpagemod_max_memsz</code>	Yes	No

**Note** The `volpagemod_max_memsz` is a VxVM tunable that is used to specify the amount of memory, measured in kilobytes, required for caching FastResync and cache object metadata. For more information on using this tunable, refer to “Performance Tuning and Monitoring” in the *VERITAS Volume Manager Administrator’s Guide*.

The tunables can be tuned either using the `vxtune` utility or by using the system-specific interface. Some of the tunables can be tuned using only the system-specific method, while the others such as the memory tunables can be tuned using both the methods. The advantage of using the `vxtune` utility to tune the parameters is that you do not require to reboot the system after modifying the tunable value. This is especially useful if you want to experiment with different values to arrive at an optimum value to suit your requirements. However, the changes to the tunables using the `vxtune` utility are

non-persistent. To make the changes persistent you will still need to use the system-specific method. However, it will require you to reboot the system for the changes to take effect.

---

**Note** The `vol_min_lowmem_sz` tunable is auto-tunable; depending on the incoming writes VVR increases or decreases the tunable value.

---

The current values for the tunables are defined in the `/etc/vx/vxtunables` file after you have used the `vxtune` utility for the first time.

### Points to note when changing the value of the tunables

- ◆ When decreasing the value of the `vol_rvio_maxpool_sz` tunable, all the RVGs on the host must be stopped.
- ◆ When decreasing the size of the tunables `vol_max_rdback_sz` and `vol_max_nmpool_sz`, pause the RLINKs.

## Changing the Tunable Values Using `vxtune`

You can use `vxtune` to display, set, or change the memory tunables that are used by VVR. The advantage of using the `vxtune` utility to tune the parameters is that you do not require to reboot the system after modifying the tunable value.

The `vxtune` utility enables you to modify the values of the following memory tunables:

```
vol_rvio_maxpool_sz
vol_min_lowmem_sz
vol_max_rdback_sz
vol_max_nmpool_sz
```

These tunable values are then updated in the `/etc/vx/vxtunables` file.

The `vxtune` utility also allows you to enable auto-tuning for the `vol_min_lowmem_sz` tunable. To do this, you must set this tunable to `-1`. However, if you do not want it to be auto-tuned then you must set it to a required value.



### ▼ To display the tunables that vxtune supports

Issue the following command on the required system to display the tunables that VVR supports along with their current values and brief description:

```
# vxtune
```

The output of this command resembles:

Tunable	Value	Description
-----	-----	-----
vol_rvio_maxpool_sz	4096	RVIO Pool Size (KBytes)
vol_min_lowmem_sz	520	Low Memory Threshold (KBytes)
vol_max_rdback_sz	4096	Readback Pool Size (KBytes)
vol_max_nmpool_sz	4096	NMCOM Pool Size (KBytes)
volpagemod_max_memsz	6144	Cache Size for FMR Metadata (KBytes)

Note that the value of the tunables is displayed in Kilobytes (KB). To display the output in bytes, use the `vxtune` command with the `-r` option as follows.

```
# vxtune -r
```

The output resembles:

Tunable	Value	Description
-----	-----	-----
vol_rvio_maxpool_sz	4194304	RVIO Pool Size (Bytes)
vol_min_lowmem_sz	532480	Low Memory Threshold (Bytes)
vol_max_rdback_sz	4194304	Readback Pool Size (Bytes)
vol_max_nmpool_sz	4194304	NMCOM Pool Size (Bytes)
volpagemod_max_memsz	6291456	Cache Size for FMR Metadata (Bytes)

### ▼ To display the values of specific tunables

Use the following command on the required host to display the value of the specific tunable:

```
# vxtune tunable_name
```

This command displays the value of the specified tunable in Kilobytes (KB).

To display the output in bytes, use the `vxtune` command with the `-r` option.

```
# vxtune -r tunable_name
```

For example, to view the value of the tunable `vol_rvio_maxpool_sz`, use the following command:

```
# vxtune vol_rvio_maxpool_sz
```

The output resembles:

```
# 65536 (KBytes)
```

### ▼ To modify the values of the tunables

Use the following command on the required host to modify the value of the specific tunable:

```
# vxtune tunable_name value
```

The value *value* of the tunable can be specified in units of KB, MB, or GB. However, the value is always displayed in kilobytes (K) regardless of the units that you have specified it in. For example, if you specify a value 500MB the value will be automatically converted and will be displayed as 512000 Kilobytes. You can use the command with the `-r` option to display the value in Bytes.

For example, to change the default value of the `vol_rvio_maxpool_sz` tunable to 128MB, use the following command.

```
# vxtune vol_rvio_maxpool_sz 128M
```

To view the changed value for the tunable use the following command

```
# vxtune vol_rvio_maxpool_sz
```

---

**Caution** Do not edit the tunable values directly in the `vxtunables` file because these changes will be ignored.

---

## Changing the VVR Tunables Using SAM

Several VVR parameters can be tuned by using the System Administration Utility (SAM). For an explanation of VVR tunables, see the *VERITAS Volume Replicator Configuration Notes*. To change the value of tunables persistently use the SAM utility. Note that the changes to the tunables will be effected only after the system has been rebooted.

### ▼ To change the value of a tunable:

1. Type `sam` to bring up the SAM interface.
2. Use the Tab key to move the control to the SAM areas display.
3. Select the `Kernel Configuration` area to display a list of options from which you must select `Configurable Parameters`.



4. Scroll to the required parameter and select it. Use the `Modify Configurable Parameter` from the `Actions` option to modify the parameter as follows:

**`tunable_name=value;`**

The change will take effect only after the next system reboot.

### **Example 1:**

To change to the modifiable tunables by using the SAM utility choose the `Configurable Parameters` option from `Kernel Configuration`.

After modifying the tunables, rebuild the kernel and reboot the system.

This chapter explains how to use VERITAS Volume Replicator (VVR) for off-host processing. The In-Band Control (IBC) Messaging feature with the Snapshot feature of VVR and the optional FastResync (FR) feature of VERITAS Volume Manager (VxVM) enable you to perform off-host processing.

This chapter explains how to perform off-host processing operations using the `vradm inbc` command. For instructions on using the `vxibc` commands to perform off-host processing operations, see Appendix C, “[Using the In-Band Control Messaging Utility vxibc and the IBC Programming API](#)” on page 371.

## What is Off-Host Processing?

Off-host processing consists of performing operations on application data on a host other than the one where the application is running. Typical operations include Decision Support Systems (DSS) and backup. In a VVR environment, off-host processing operations can be performed on the Secondary of the Replicated Data Set. This reduces the load on the application server, the Primary.

The model for data access on the Secondary is that you break off a mirror from each data volume in the RVG, perform the operation on the mirror, and then reattach the mirror while replication is in progress.

## In-Band Control Messaging Overview

When you take a snapshot on the Secondary, it contains a point-in-time copy of the data on the Primary. Because the Secondary may be behind the Primary, it is not known at exactly what time this point-in-time copy was made.

IBC messaging enables you to send a message in the replication stream to notify the Secondary that an event has occurred on the Primary. In the case of a file system, you can use the `sync` command on the Primary, and then send an IBC message. When this message arrives on the Secondary, the data on the Secondary is consistent at the file system level and replication stops. You then split off a mirror, which now contains a consistent image of the file system, and unfreeze replication.



The model with IBC Messaging is that a process on the Secondary waits for the IBC Message, and a process on the Primary sends the message when the desired event has occurred.

VVR provides the following options for IBC Messaging:

- ◆ Single command to perform off-host processing operations—`vradmin ibc`

This chapter explains the function of IBC Messaging and how to use the command `vradmin ibc` for off-host processing.

- ◆ IBC Messaging command-line utility—The `vxibc` utility

For more information on how to use the command-line utility `vxibc`, see [“Using the In-Band Control Messaging Utility `vxibc` and the IBC Programming API”](#) on page 371.

---

**Note** If you choose not to use IBC Messaging, the data on the Secondary is consistent and can be recovered by the application but it is out-of-date.

---

## How to Use the Data on the Secondary

### Using Snapshots

A snapshot is an image of the online data volumes at a specific point-in-time. Use snapshots of the Secondary data volumes to perform off-host processing operations, instead of directly using the Secondary data volumes. The data on the original volumes may change but the data on the snapshot can still be used as a stable and independent copy for various purposes.

VVR provides two methods of creating snapshots: instant snapshots and traditional snapshots. The instant snapshot feature is a separately licensed feature of VVR. For more information on instant snapshots, see [“Using the Instant Snapshot Feature”](#) on page 157. VVR also provides sample IBC scripts that can be used for creating snapshots. For more information, see [“Sample Scripts”](#) on page 210.

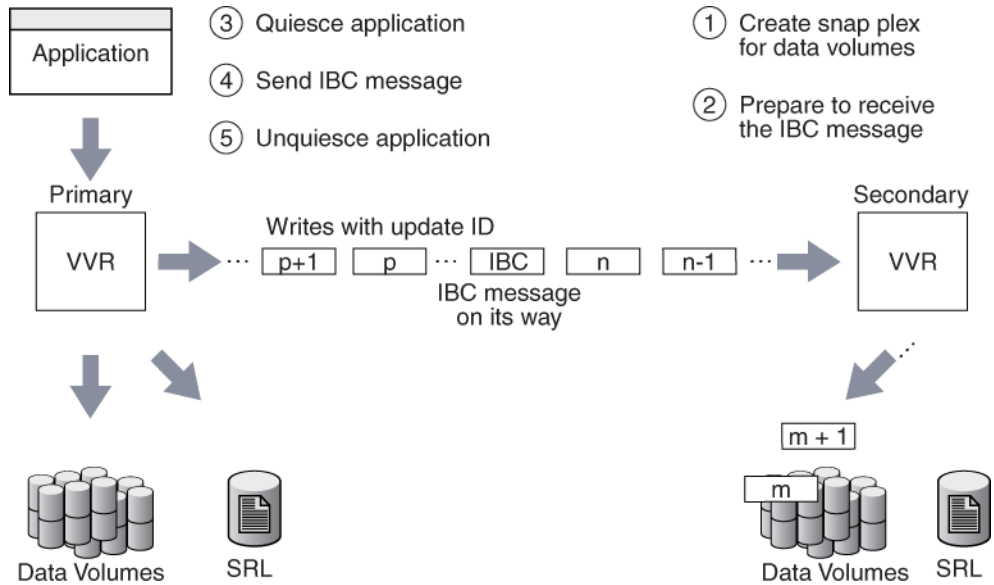
With the traditional snapshot feature, depending on the size of the volume, the time required for initial synchronization of the plexes can be very large. For more information on traditional snapshots, see [“Using the Traditional Snapshot Feature”](#) on page 177.

Before you can use the snapshot, some application-dependent recovery procedure has to be performed. For example, if the volume contains a file system, run the `fsck` program before mounting the file system.



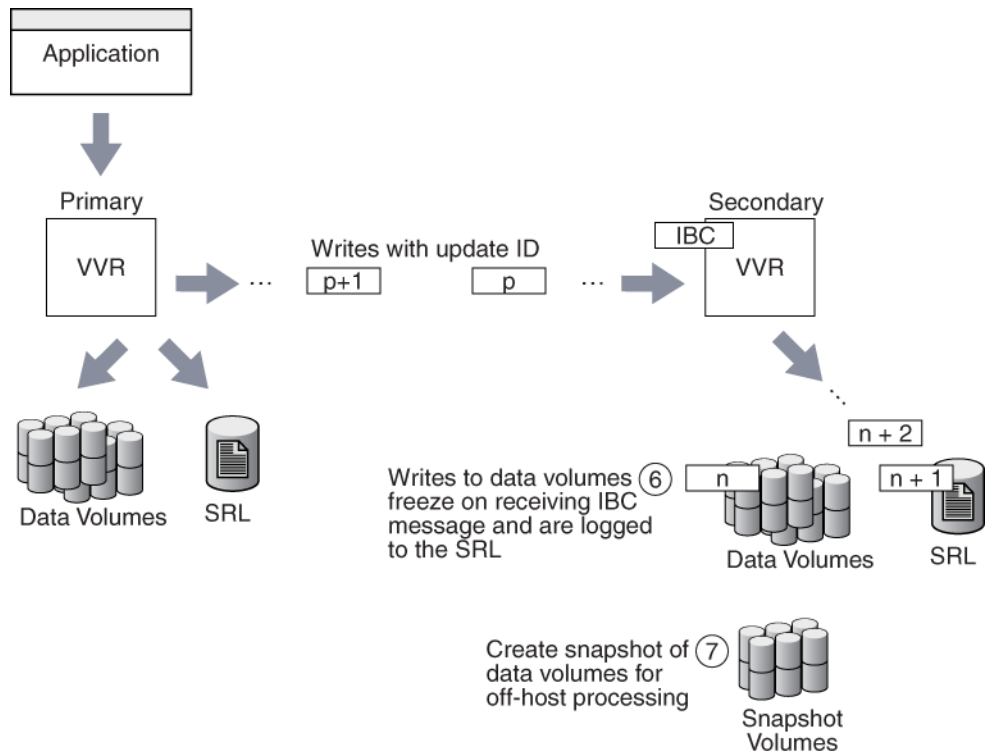
# In-Band Control Messaging Explained

You can use IBC Messaging to notify the Secondary that the data volumes in the Primary RVG are consistent at the application-level.



Using IBC messaging, you can inject a user-defined control message into the update stream of an RVG at a point when the Primary is consistent at the application level. When the IBC message reaches the Secondary, data volumes on the Secondary are frozen and any new updates received after the IBC message are logged into the Secondary SRL. As a result, the Secondary does not reflect further updates to data volumes until the user acknowledges the IBC message.

The Secondary data volumes are now consistent at the application level.



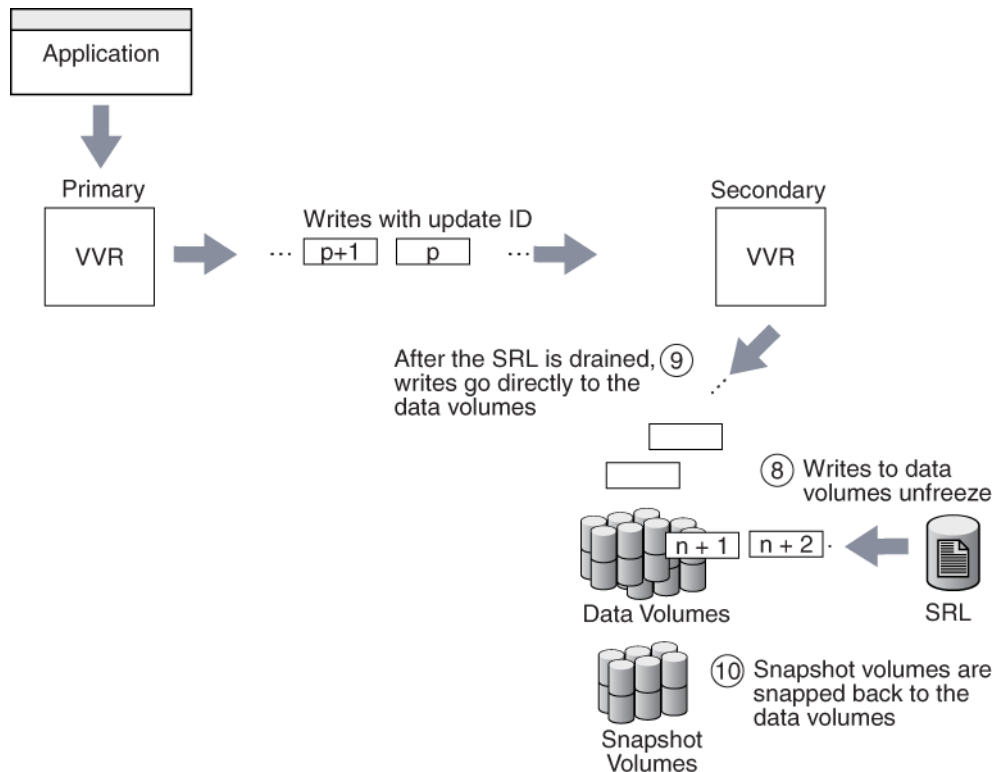
While replication is frozen, take snapshots of the data volumes on the Secondary. The created snapshot volumes are consistent at the application level and require less time to recover when the application on the Secondary starts.

An application must be in a quiesced mode before the IBC message is sent to the Secondary, to achieve application-level consistency. For a database application running on the Primary host, an IBC message can be inserted at a point at which the application considers its raw volume contents to be consistent, such as during the database *hot-backup* mode.

In the case of a file system, when you enter the `sync` command on the Primary to flush the previously unwritten file system buffers to the data volume, the file modifications up to that point are saved to the data volume. You can insert an IBC message on the Primary to notify the Secondary that the `sync` command is complete. In general, there is no way to keep the file system in the synced state while you generate the IBC; however, if this is done as soon as the sync completes, there should be little to recover on the Secondary.

Even if you are using synchronous replication, IBC Messaging is useful to notify the Secondary when the consistency point is reached and to ensure that the writes on the Secondary are written to the data volumes before the snapshot is taken.

When the IBC reaches the Secondary, the subsequent updates are logged into the SRL, and the data volumes are frozen. Now, the Secondary data volumes are consistent at the application level, and you can take a snapshot. If you take a backup of the snapshot volumes, the file systems on the backup volumes are consistent.



An IBC message is delivered to the Secondary host in causal order with respect to other activity on the data volumes. Before delivery of the message to the Secondary host, any previous update activity is flushed. You have the option of allowing subsequent updates to be applied immediately to the Secondary data volumes, or logging the subsequent updates into the SRL until released by unfreezing the Secondary RVG. Since the purpose of the IBC message is to achieve a sync point on the Secondary, choose the option to stop subsequent updates until the snapshot is taken. This is the default option in the examples below.

IBC messages are guaranteed to be delivered at least once and might be delivered more than once if an error such as a network outage or machine crash occurs during the delivery. The script you write to perform the IBC operation must be able to manage receiving the multiple delivery of the same IBC message.

## Performing Off-Host Processing Tasks

VVR enables you to integrate application preparation tasks, IBC messaging, and off-host processing tasks using the single command `vradmin ibc`. The `vradmin ibc` command simplifies the off-host processing task by providing scripts for application-specific tasks. The `vradmin ibc` command uses IBC Messaging and executes a set of user-defined scripts to perform the required off-host processing task. You are not required to remember the sequence in which to perform these tasks. The `vradmin ibc` command simplifies the off-host processing process by performing the following operations:

- ◆ Executes application-specific scripts and the scripts to be used for off-host processing in the required sequence.
- ◆ Executes these scripts on the appropriate host, that is, Primary or the Secondary.
- ◆ Inserts the IBC message at the appropriate time.

## Tasks to Perform for Off-Host Processing

Following is the typical sequence of tasks that must be performed to use IBC Messaging for off-host processing:

1. Prepare the Secondary for off-host processing. For example, create snapshot plexes on the data volumes on the Secondary (`prefreeze` task).
2. Register an application name for sending and receiving IBC messages on the Primary and the Secondary. Prepare to receive the IBC message on the Secondary.
3. Quiesce the application on the Primary (`quiesce` task).
4. Send the IBC message from the Primary to the Secondary.
5. Unquiesce the application on the Primary (`unquiesce` task).
6. Perform the off-host processing task on the Secondary after the Secondary receives the IBC message and replication freezes (`onfreeze` task). Note that the updates to the Secondary data volume is frozen while the off-host processing task is in progress.
7. Unfreeze the Secondary after the off-host processing task completes.
8. Perform additional tasks on the Secondary after replication has resumed. For example, reattaching the snapshot volumes to the original data volumes on the Secondary (`postfreeze` task).
9. Unregister the application on both the Primary and the Secondary.

The single command `vradmin ibc` can be used to perform this sequence of tasks. To use the `vradmin ibc` command, you need to provide scripts named `prefreeze`, `quiesce`, `unquiesce`, `onfreeze`, `postfreeze` to perform the tasks in step 1, step 3, step 5, step 6, and step 8 respectively. The `vradmin ibc` command uses these scripts with IBC Messaging to perform these tasks in sequence. For more information about the user-defined scripts, see [“Understanding the Scripts Used for the vradmin ibc Command”](#) on page 208.

You can also use the `vxibc` commands to perform off-host processing operations. For details on using the available `vxibc` commands, see [“Using the In-Band Control Messaging Utility vxibc and the IBC Programming API”](#) on page 371.



## Using the IBC Messaging Command `vradmin ibc`

The `vradmin ibc` command enables you to perform off-host processing operation in a single command.

### Prerequisites for Using `vradmin ibc` Command

- ✓ The Primary RLINK that points to the Secondary host participating in the `vradmin ibc` command must be in the `CONNECT` state.
- ✓ The `onfreeze` script must exist on each Secondary host participating in the `vradmin ibc` command.
- ✓ Make sure each user-defined script to be used in the `vradmin ibc` command exits with a status of 0 on successful completion and a status of nonzero on unsuccessful completion.
- ✓ The user-defined scripts must have `execute` permissions for `root` user.

---

**Caution** The `vradmin ibc` executes scripts using root privileges. If the scripts can be modified by a non-privileged user, there is a potential security risk. To prevent this, ensure that you have the proper access privileges set on the scripts used with the `vradmin ibc` command.

---

### ▼ To perform an off-host processing task on one or more Secondary RVGs in an RDS

1. Make sure the RLINKs are in the `CONNECT` state. If the RLINKs are not in the `CONNECT` state, use the `vradmin startrep` command to start replication.
2. Create the user-defined scripts for this off-host processing task.
  - a. Create the following directory on all hosts participating in the `vradmin ibc` command:

```
/etc/vx/vvr/ibc_scripts/task_name
```

where `task_name` is the name of the off-host processing task and is the same as the `task_name` argument used in the `vradmin ibc` command.
  - b. Create the appropriate scripts for the required off-host processing task and copy the scripts in the directory created in [step a](#). See [“Understanding the Scripts Used for the `vradmin ibc` Command”](#) on page 208.

3. Run the following command from any host in the RDS:

```
# vradmin -g diskgroup ibc rvg_name task_name [sechost]...
```

The argument *diskgroup* represents the disk group that contains the RVG on the local host.

The argument *rvg\_name* is the name of the RVG on the local host and represents its RDS.

The argument *task\_name* is the name of the off-host processing task and is the same as the name of the directory created in [step a](#).

The argument *sechost* is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command. The argument *sechost* is optional if the RDS contains only one Secondary. To perform the task on multiple Secondary hosts, specify a space-separated list with the name of each Secondary to be included. Use the `-all` option to perform the task on all the Secondary hosts in the RDS.

### Example—Creating a Snapshot on the Secondary Using the `vradmin ibc` Command

This example shows how to create a snapshot of the data volumes on the Secondary `london` using the `vradmin ibc` command. The RVG `hr_rvg`, which belongs to the disk group `hrdg`, has been created on the Primary and Secondary. This example also assumes that Secondary data volumes have associated snapshot plexes. It uses the application name `dss_app`.

1. Create the following directory on Secondary host:

```
/etc/vx/vvr/ibc_scripts/dss_app
```

2. Create the `onfreeze` script in the `/etc/vx/vvr/ibc_scripts/dss_app` directory on the Secondary host by including the following command to create the snapshot of the data volumes on the Secondary:

```
#!/bin/sh
/usr/sbin/vxrvrg -g hrdg snapshot hr_rvg
```

3. On the Primary, put the application using the Primary data volumes in the quiesce mode.
4. Create the snapshot by running the following command on any host in the RDS:

```
# vradmin -g hrdg ibc hr_rvg dss_app london
```

5. On the Primary, take the application out of quiesced mode to resume it.



## Understanding the Scripts Used for the vradmin ibc Command

The `vradmin ibc` command executes the user-defined scripts: `prefreeze`, `quiesce`, `unquiesce`, `onfreeze`, and `postfreeze`. Note that the `onfreeze` script is mandatory and must be present on the Secondary. The scripts `prefreeze`, `quiesce`, `unquiesce`, and `postfreeze` are optional. However, if you provide the `quiesce` script, you must provide the `unquiesce` script and vice versa. You must name the user-defined scripts `prefreeze`, `quiesce`, `unquiesce`, `onfreeze`, or `postfreeze`.

---

**Note** A user-defined script can either be a shell script or a binary.

---

### Location of the scripts

The scripts must reside in the `/etc/vx/vvr/ibc_scripts/task_name` directory on the Primary and the Secondary host. Note that `task_name` is the name of the off-host processing task and is the same as the `task_name` argument used in the `vradmin ibc` command. For example, if the off-host processing task is Decision Support Systems (DSS), you can choose a task name of `dss`; or if the off-host processing task is Backup, you can choose the task name of `backup`.

---

Host	Directory	Name of Script
Primary	<code>/etc/vx/vvr/ibc_scripts/task_name</code>	<code>quiesce</code>
Primary	<code>/etc/vx/vvr/ibc_scripts/task_name</code>	<code>unquiesce</code>
Secondary	<code>/etc/vx/vvr/ibc_scripts/task_name</code>	<code>onfreeze</code>
Secondary	<code>/etc/vx/vvr/ibc_scripts/task_name</code>	<code>prefreeze</code>
Secondary	<code>/etc/vx/vvr/ibc_scripts/task_name</code>	<code>postfreeze</code>

---



When the `vradmin ibc` command executes each script, it passes the following arguments to the script:

Parameter	Description
Argument 1	Name of the disk group of the Primary RVG.
Argument 2	Name of the Primary RVG.
Argument 3	The <code>task_name</code> specified in the <code>vradmin ibc</code> command.
Remaining Arguments	Names of the RLINKs participating in the <code>vradmin ibc</code> command.

The following section describes how each script is used with the `vradmin ibc` command:

- ◆ The `prefreeze` script

Use this script on the Secondary to prepare for the tasks to be performed in the `onfreeze` script while the replication to Secondary is frozen. For example, if you want to take a snapshot of the Secondary data volumes while the replication on the Secondary is frozen, the `prefreeze` script can be used to add snapshot plexes to the Secondary data volumes to prepare for the `snapshot` command.

- ◆ The `quiesce` script

The `vradmin ibc` command executes the `quiesce` script on the Primary before it sends the IBC message to the Secondary. Use this script to quiesce the application running on the Primary RVG and to make the Primary data volumes consistent at the application level. The `vradmin ibc` command injects an IBC message in a small amount of time, and hence the duration for which the application remains quiesced is small.

- ◆ The `unquiesce` script

The `vradmin ibc` command executes this script on the Primary after it sends the IBC message to the Secondary. Use this script to resume the application running on the Primary if the application was quiesced.

- ◆ The `onfreeze` script

The `vradmin ibc` command executes this script on the Secondary while replication on the Secondary is frozen after receiving the IBC message from the Primary. Use this script to perform the required off-host processing operation, for example, taking a snapshot of the Secondary data volumes.



- ◆ The `postfreeze` script

The `vradm in ibc` command executes this script on the Secondary after it executes the `onfreeze` script and after the replication on the Secondary is unfrozen. For example, if a snapshot of the Secondary data volumes was taken in the `onfreeze` script, this script can be used to reattach the snapshot volumes to the Secondary data volumes.

## Sample Scripts

The `/etc/vx/vvr/ibc_scripts` directory contains the following sample script directories:

```
sample_db_snapshot
sample_vxfs_snapshot
sample_so_snapshot
```

These sample scripts show how to use the user-defined scripts with the `vradm in ibc` command. Refer to the *README* file provided in `/etc/vx/vvr/ibc_scripts` directory for instructions on how to use the sample scripts to perform off-host processing tasks.

---

**Note** Sample scripts are provided for reference. Customize the sample scripts to suit your requirements.

---

## Examples of Off-Host Processing

The examples in this chapter assume that the following VVR configuration has been set up on the Primary and Secondary hosts:

Name of the Primary host: `seattle`

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

Name of the Secondary host: `london`

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

These examples use the application name `dss_app` for off-host processing tasks.



## Example 1—Decision Support Using the Snapshot Feature and the `vradm ibc` Command

1. Create the following directory on both Primary and Secondary hosts:

```
/etc/vx/vvr/ibc_scripts/dss_app
```

2. Create the `quiesce` and `unquiesce` scripts and copy them to the `/etc/vx/vvr/ibc_scripts/dss_app` directory on the Primary host.

In the `quiesce` script, provide commands to put the application that is using the Primary data volumes `hr_dv01` and `hr_dv02` in quiesce mode.

In the `unquiesce` script, provide commands to resume the application or take it out of the quiesce mode.

3. Create the `prefreeze` and `onfreeze` scripts and copy them to the `/etc/vx/vvr/ibc_scripts/dss_app` directory on the Secondary host.

In the `prefreeze` script, include the following commands to add snapshot plexes to the Secondary data volumes `hr_dv01` and `hr_dv02`:

```
#!/bin/sh  
  
/usr/sbin/vxassist -g hrdg snapstart hr_dv01  
  
/usr/sbin/vxassist -g hrdg snapstart hr_dv02
```

In the `onfreeze` script, include the following command to take the snapshot of the Secondary volumes:

```
/usr/sbin/vxrvrg -g hrdg snapshot hr_rvg
```

---

**Note** This example does not need a `postfreeze` script.

---

4. Run the following `vradm ibc` command from any host in the RDS:

```
# vradm -g hrdg ibc hr_rvg dss_app london
```

5. On the Secondary, use the snapshot data volumes `SNAP-hr_dv01` and `SNAP-hr_dv02` to run the DSS application, that is, for off-host processing.
6. When the DSS application completes, reattach the snapshot plexes to the data volumes by issuing the following command on the Secondary host `london`:

```
# vxrvrg -g hrdg snapback hr_rvg
```

The `snapback` destroys the `SNAP` volumes and reattaches the snapshot plexes to their original volumes. If you have enabled `FR` on these volumes, the reattach is faster.

## Example 2—Backing Up Using the Snapshot Feature and the `vradm in ibc` Command

1. Perform [step 1](#) and [step 2](#) from “[Example 1—Decision Support Using the Snapshot Feature and the `vradm in ibc` Command](#)” on page 212.
2. Create the `prefreeze`, `onfreeze`, and `postfreeze` scripts and copy them in the `/etc/vx/vvr/ibc_scripts/dss_app` directory on the Secondary host.

In the `prefreeze` script, include the following commands to add snapshot plexes to the Secondary data volumes `hr_dv01` and `hr_dv02`:

```
#!/bin/sh
/usr/sbin/vxassist -g hrdg snapstart hr_dv01
/usr/sbin/vxassist -g hrdg snapstart hr_dv02
```

In the `onfreeze` script, include the following command to take the snapshot of the Secondary volumes:

```
/usr/sbin/vxrvrg -g hrdg snapshot hr_rvg
```

In the `postfreeze` script, include the following commands to back up the snapshot volumes and to reattach the snapshot plexes to the original volumes after the backup is complete:

```
# Back up data from SNAP-hr_dv01 and SNAP-hr_dv02
dd if=/dev/vx/rdisk/hrdg/SNAP-hr_dv01 of=/dev/rmt/0
dd if=/dev/vx/rdisk/hrdg/SNAP-hr_dv02 of=/dev/rmt/0
# Reattach the snapshot volumes to the Secondary data volumes
/usr/sbin/vxrvrg -g hrdg snapback hr_rvg
```

3. Run the following `vradm in ibc` command from any host in the RDS:

```
# vradm in -g hrdg ibc hr_rvg dss_app london
```



## Example 3—Performing Block-Level Backup of the Secondary Data Using the `vradmin ibc` Command

This method performs the backup directly from the Secondary data volumes while the replication to Secondary is frozen. Make sure the Secondary SRL is large enough to hold the writes sent from the Primary while the backup is in progress. In this method, the Secondary data volumes are under replication control, and hence you cannot write to them. Use this method only if the backup process does not perform any writes to the Secondary data volumes.

1. Perform [step 1](#) and [step 2](#) from “[Example 1—Decision Support Using the Snapshot Feature and the `vradmin ibc` Command](#)” on page 212.
2. Create the `onfreeze` script and copy it in the `/etc/vx/vvr/ibc_scripts/dss_app` directory on Secondary host:

In the `onfreeze` script, include the following commands to perform block-level backup of the Secondary data volumes:

```
#!/bin/sh
dd if=/dev/vx/rdisk/hrdg/hr_dv01 of=/dev/rmt/0
dd if=/dev/vx/rdisk/hrdg/hr_dv02 of=/dev/rmt/0
```

---

**Note** This example does not need the `prefreeze` and `postfreeze` scripts.

---

3. Run the following `vradmin ibc` command from any host in the RDS:

```
# vradmin -g hrdg ibc hr_rvg dss_app london
```

## Transferring the Primary Role

---

In a VVR environment, applications can only write to the Primary data volumes. When replication is active, the application cannot write to the Secondary data volumes. To start the applications on the Secondary, you must transfer the Primary role to the Secondary. After transferring the role, you can start the application on the new Primary.

VVR enables you to transfer the Primary role from a healthy or failed Primary using a single command. It also enables you to fail back to the original Primary using a simple set of commands.

VVR offers the following methods to transfer the Primary role:

- ◆ [Migrating the Primary](#)
- ◆ [Taking Over from an Original Primary](#)
- ◆ [Failing Back to the Original Primary](#)

This chapter explains how to transfer the Primary role using each of these methods.

---

**Note** If the RDS is configured in a VERITAS Cluster Server (VCS) environment, use the `hagrp` command to offline and online the corresponding resources. For more information on offlining and onlining resources, see the *VERITAS Cluster Server User's Guide*. If you are using the RVGPrimary agent for VVR to manage role transfer, refer to the *VCS Agents for VVR Configuration Guide*.

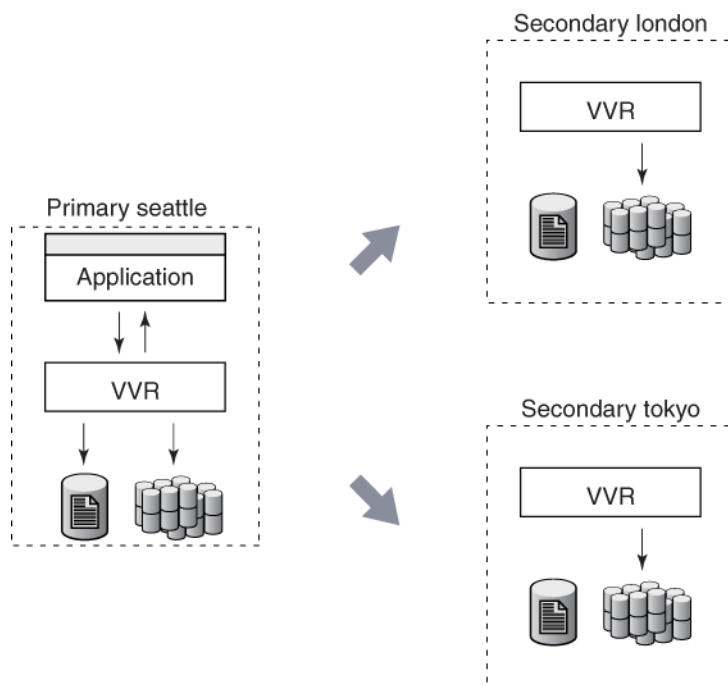
---



## Migrating the Primary

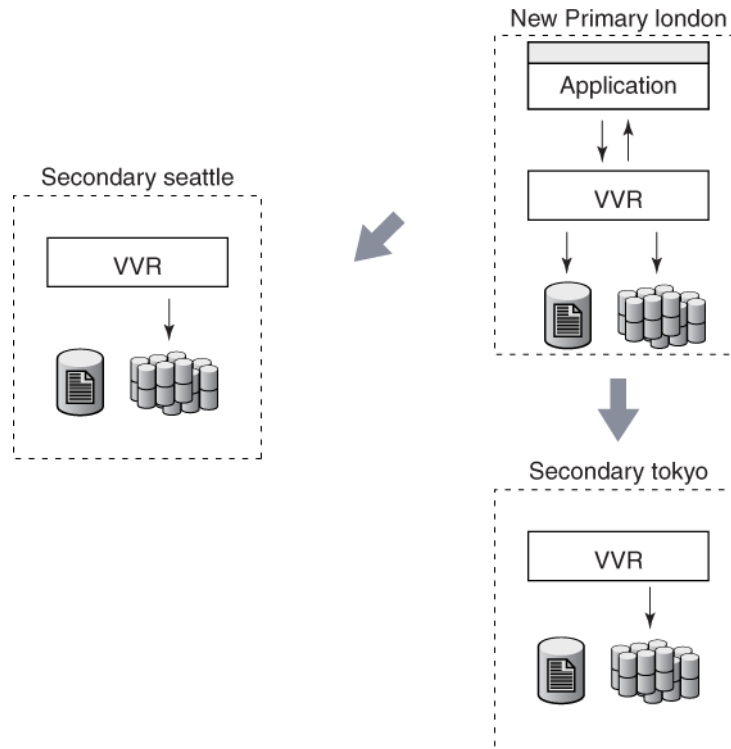
Migration involves transferring a healthy Primary of a Replicated Data Set (RDS) to a Secondary when the application involved in replication is inactive. Migration of a Primary to a Secondary is useful when the Primary must be brought down for reasons such as maintenance or to make the application active on another node.

In the following illustration, the Primary `seattle` is replicating to the Secondary hosts `london` and `tokyo`.





In the following illustration, the Primary role has been migrated from `seattle` to `london` and the new Primary `london` is replicating to `seattle` and `tokyo`. If you had already created RLINKs between `london` and `tokyo` when setting up replication, you do not need to manually reconfigure the additional Secondary `tokyo` as a part of the RDS. It will automatically be added as a Secondary of the new Primary `london`.



VVR provides the `vradmin migrate` command to migrate a healthy Primary. The `vradmin migrate` command performs the following functions:

- ◆ Migrates the Primary role of an RDS to a Secondary, thus converting the Secondary RVG to the Primary RVG.
- ◆ Converts the original Primary of the RDS to Secondary in the RDS.
- ◆ Reconfigures the original Primary and the new Primary.

The `vradmin migrate` command restores the original configuration if it fails on any of the hosts during its execution.



Before migrating the Primary role, the `vradmin migrate` command displays a warning and prompts the user to confirm whether or not all the applications using the Primary volumes are stopped. To skip this confirmation, use the `-s` option with the `vradmin migrate` command, which proves useful in scripts.

### Prerequisites for migrating the Primary:

- ✓ The data volumes in the RDS must be inactive, that is, applications that use the Primary data volumes must be stopped.
- ✓ All Secondaries must be up-to-date.
- ✓ All attached RLINKs must be in the `CONNECT` state.

To migrate a healthy Primary from any host in an RDS:

```
# vradmin -g diskgroup migrate local_rvgname newprimary_name
```

The argument `diskgroup` is the disk group on the local host

The argument `local_rvgname` is the name of the RVG on the local host

The argument `newprimary_name` is the name of the new Primary host, that is, the existing Secondary host. For an RDS with only one Secondary, this argument is optional. Note that the `newprimary_name` argument must be the host name displayed by the `vradmin printrvg` command.

## Important Notes for Migrating the Primary Role

- ◆ We recommend that you set the size of the SRL the same on the Primary and Secondary nodes because any of the Secondary nodes could be later converted to a Primary by using the `vradmin migrate` or the `vradmin takeover` command.
- ◆ We recommend that you configure the Primary and Secondary data volumes with the same names. However, if the names of the data volumes on the Primary and the Secondary do not match, map the names of the Secondary data volumes with the differently named Primary data volumes. For more information, see [“Mapping the Name of a Secondary Data Volume to a Differently Named Primary Data Volume”](#) on page 401.
- ◆ For an RDS with multiple Secondaries:
  - ◆ We recommend that you wait until all the Secondaries are up-to-date before migrating the Primary role. The `vradmin migrate` command fails if all the Secondaries are not up-to-date. If the additional Secondaries were out of date before performing the migrate operation, then you would need to perform a complete synchronization.

- ◆ After successful migration, if you had already created RLINKs between every pair of secondaries while setting up replication, you do not need to manually reconfigure the additional secondaries as a part of the RDS. Otherwise, you must reconfigure them manually. For more information, see [“Example 2—Migrating the Primary Role in a Setup with Multiple Secondaries”](#) on page 220.

## Example 1—Migrating from a Healthy Primary

This example explains how to migrate the original Primary `seattle` to the Secondary host `london`.

---

**Note** Create SRLs of the same size on the Primary and Secondary hosts.

---

Before migration, the configuration of the RDS looks like this:

	On Primary	On Secondary
<b>Host Name</b> displayed by the <code>vradm</code> <code>printrvg</code> command	<code>seattle</code>	<code>london</code>
<b>RVG</b>	<code>hr_rvg</code>	<code>hr_rvg</code>
<b>RLINK</b>	<code>rlk_london_hr_rvg</code>	<code>rlk_seattle_hr_rvg</code>

After the migration, the configuration of the RDS looks like this:

	On Primary	On Secondary
<b>Host Name</b> displayed by the <code>vradm</code> <code>printrvg</code> command	<code>london</code>	<code>seattle</code>
<b>RVG</b>	<code>hr_rvg</code>	<code>hr_rvg</code>
<b>RLINK</b>	<code>rlk_seattle_hr_rvg</code>	<code>rlk_london_hr_rvg</code>

### ▼ To migrate the Primary RVG `hr_rvg` to host `london`:

1. Stop the applications that use the Primary data volumes. For example, if the application is a file system, unmount it.
2. Verify that the Primary RLINK is up-to-date by using the `vxrlink status` command. On the Primary `seattle`, issue the following command:

```
# vxrlink -g hrdg status rlk_london_hr_rvg
```



The `vradmin migrate` command fails if the Primary RLINK is not up-to-date or not in the `CONNECT` state. It also fails if the data volumes are active.

3. Migrate the Primary RVG `hr_rvg` by typing the following command from any host in the RDS:

```
# vradmin -g hrdg migrate hr_rvg london
```

`london` is the name of the Secondary host displayed by the `vradmin printrvg` command.

4. Restart the application.

Because the application was stopped properly before the migration, an application recovery is not required.

By default, the `vradmin migrate` command enables replication from the new Primary `london`. To start the application before enabling replication, first, issue the `vradmin pauserrep` command, start the application, and then resume replication.

## Example 2—Migrating the Primary Role in a Setup with Multiple Secondaries

We recommend that you create RLINKs between hosts `london` and `tokyo` when setting up the RDS.

---

**Note** Create SRLs of the same size on the Primary and Secondary hosts.

---

Before migration, the configuration of the RDS looks like this:

	On Primary	On Secondary	On Secondary
<b>Host Name</b> displayed by the command: <code>vradmin printrvg</code>	seattle	london	tokyo
<b>RVG</b>	hr_rvg	hr_rvg	hr_rvg
<b>RLINKs</b>	rlk_london_hr_rvg*	rlk_seattle_hr_rvg*	rlk_seattle_hr_rvg*
	rlk_tokyo_hr_rvg*	rlk_tokyo_hr_rvg	rlk_london_hr_rvg

\* Indicates active RLINKs

---

▼ **To migrate the Primary RVG `hr_rvg` to host `london`:**

1. Stop the applications that use the Primary data volumes. For example, if the application is a file system, unmount it.
2. Verify that the Primary RLINKs are up-to-date by using the `vxrlink status` command. On the Primary `seattle`, issue the following commands:

```
# vxrlink -g hrdg status rlk_london_hr_rvg
# vxrlink -g hrdg status rlk_tokyo_hr_rvg
```

The `vradm migrate` command fails if the Primary RLINKs are not up-to-date or not in the `CONNECT` state. It also fails if the data volumes are active.

3. Migrate the Primary RVG `hr_rvg` by typing the following command from any host in the RDS:

```
# vradm -g hrdg migrate hr_rvg london
```

4. If you had created RLINKs between the Secondary `london` and the additional Secondary `tokyo`, host `tokyo` is automatically added to the new configuration.

Otherwise, you must manually add `tokyo` as a Secondary to the new Primary `london`. To do this, create RLINKs between `london` and `tokyo` and associate them to the respective RVGs using the following commands.

On host `london`:

```
# vxmake -g hrdg rlink rlk_tokyo_hr_rvg local_host=london \
  remote_host=tokyo remote_rlink=rlk_london_hr_rvg \
  remote_dg=hrdg
# vxrlink -g hrdg assoc hr_rvg rlk_tokyo_hr_rvg
```

On host `tokyo`:

```
# vxmake -g hrdg rlink rlk_london_hr_rvg local_host=tokyo \
  remote_host=london remote_rlink=rlk_tokyo_hr_rvg \
  remote_dg=hrdg
# vxrlink -g hrdg assoc hr_rvg rlk_london_hr_rvg
```

---

**Note** By default, the `vxmake rlink` command creates the RLINK with the protocol set to UDP/IP. To change the protocol to TCP/IP, see “[Setting the Network Transport Protocol](#)” on page 54.

---



5. Start replication to `tokyo` using the following command:

```
# vradmin -g hrdg -f startrep hr_rvg tokyo
```

---

**Note** Ensure that the above command is run before starting the application on the new Primary `london`.

---

6. Restart the application.

Because the application was stopped properly before the migration, an application recovery is not required.

By default, the `vradmin migrate` command enables replication from the new Primary `london`. To start the application before enabling replication, first, issue the `vradmin pauserep` command, start the application, and then resume replication.

After migration, the configuration of the RDS looks like this:

---

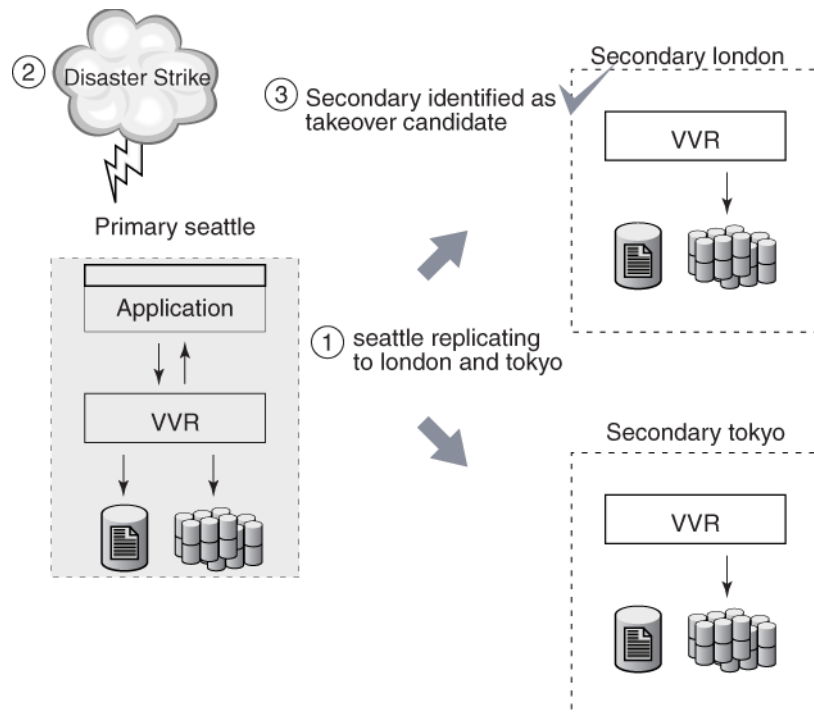
	On Primary	On Secondary	On Secondary
<b>Host Name</b>	<code>london</code>	<code>seattle</code>	<code>tokyo</code>
displayed by the command: <code>vradmin printrvg</code>			
<b>RVG</b>	<code>hr_rvg</code>	<code>hr_rvg</code>	<code>hr_rvg</code>
<b>RLINKs</b>	<code>rlk_seattle_hr_rvg*</code>	<code>rlk_london_hr_rvg*</code>	<code>rlk_london_hr_rvg*</code>
	<code>rlk_tokyo_hr_rvg*</code>	<code>rlk_tokyo_hr_rvg</code>	<code>rlk_seattle_hr_rvg</code>
* Indicates active RLINKs			

---

## Taking Over from an Original Primary

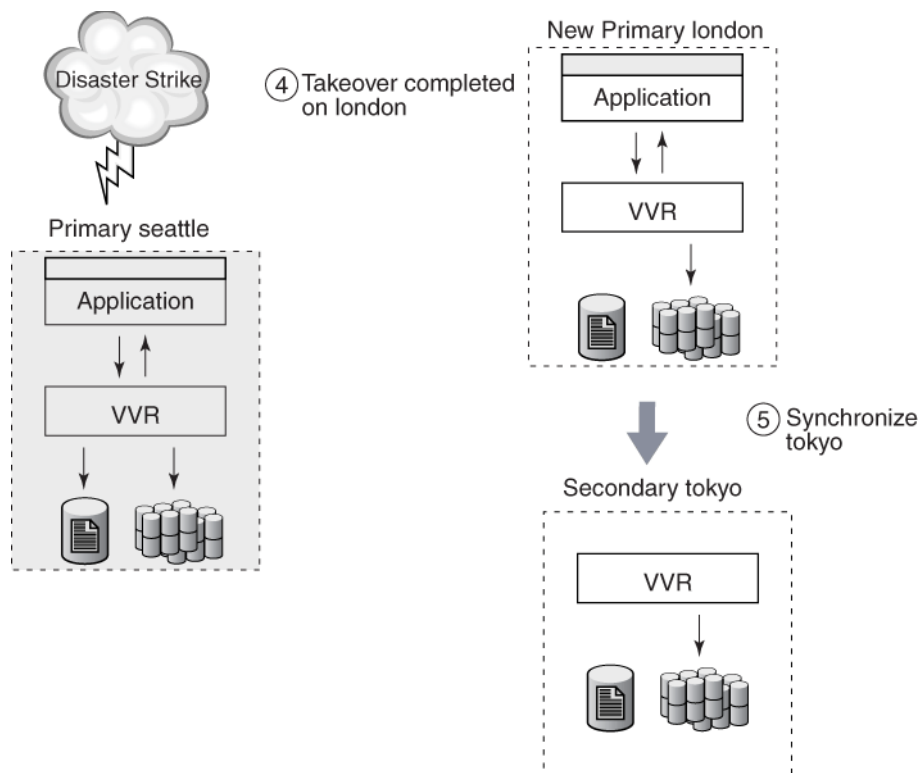
The takeover procedure involves transferring the Primary role from an original Primary to a Secondary. When the original Primary fails or is destroyed because of a disaster, the takeover procedure enables you to convert a consistent Secondary to a Primary. The takeover of a Primary role by a Secondary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

In the following illustration, the Primary *seattle* is replicating to the Secondary hosts *london* and *tokyo* when disaster strikes the Primary *seattle*. The Secondary *london* has been identified as the Secondary for takeover and will become the New Primary after the takeover is complete.



After the takeover is complete, the Secondary london becomes the new Primary. If you had already created RLINKs between london and tokyo when setting up replication, you do not need to manually reconfigure the additional Secondary tokyo as a part of the RDS. It will automatically be added as a Secondary of the new Primary london.

You must then synchronize tokyo with the new Primary london and start replication.



VVR provides the `vradmin takeover` command to transfer the Primary role from an original Primary to a Secondary. This command can be run from the Secondary host only when the Primary is not reachable from the Secondary on which the takeover command is to be run. Upon successful completion of the takeover, the Secondary becomes the Primary. VVR displays an error message if you enter the `vradmin takeover` command on the Primary.

For configurations with multiple Secondaries, you can use the `vxrlink updates` command to find the most suitable replacement for the failed Primary. For more information, see [“Identifying the Most Up-to-Date Secondary”](#) on page 114.



## Important Notes About Taking Over from an Original Primary

- ◆ The Secondary that is to become the new Primary must be consistent before taking over the Primary role. Any consistent Secondary can be selected to be the new Primary. A Secondary that is undergoing DCM resynchronization or autosync is inconsistent and cannot be used as a target of a takeover operation. Use the `vxprint -l` on the Secondary `RLINK` to check whether the `consistent` flag is set.
- ◆ Writes that were not replicated to the Secondary before the Primary became unusable are lost.

To preserve the data that was not replicated to the Secondary before the Primary became unusable, we recommend that you take a snapshot of the Primary data volumes before starting takeover with fast failback synchronization. The applications can later be started from the snapshot and you can apply the transactions or files that did not get replicated, to the active data.

- ◆ The Primary role takeover must be based on Recovery Point Objective (RPO) or the Recovery Time Objective (RTO) depending on the specific business requirement. For example, consider the scenario where the Secondary was behind by 100MB at the time of Primary crash and if the Primary is not expected to be available for another four hours, you will need to decide between having the application up and available on the Secondary or waiting for the Primary to become available after four hours. If it is required that the application must be available immediately, then the data on the Primary that was not yet replicated will be lost. Thus, takeover can result in data loss.
- ◆ The Primary role takeover is intended to support disaster recovery applications. Only a limited number of error scenarios prior to loss of the Primary node can prevent a successful takeover. These error scenarios leave the Secondary RVG in an inconsistent state and prevent takeover. All such scenarios involve a hardware failure of a Secondary data volume or Secondary SRL, and as a result, the Secondary RVG will be in an inconsistent state. The chances of such an occurrence are reduced if these volumes are configured (locally) as mirrored volumes.

---

**Note** The takeover procedure does not guarantee that the new Primary and any additional Secondary RVGs have identical contents. The remaining Secondaries must be completely synchronized with the new Primary.

---

- ◆ We recommend that you set the size of the SRL the same on the Primary and Secondary nodes because any of the Secondary nodes could be later converted to a Primary by using a `migrate` or `takeover` command.
- ◆ Each data volume on the new Primary must have a Data Change Map (DCM) associated with it.



The `vradmin takeover` command performs the following functions on the RDS to which the original Primary belongs:

- ◆ Converts the Secondary RVG to Primary RVG.
- ◆ Enables the fast failback feature on the new Primary, which speeds up failback to the original Primary when it recovers. The fast failback feature is described in the next section.

### About Fast Failback

The applications are started on the new Primary after the takeover is complete. The fast failback feature uses failback logging for incremental synchronization of the original Primary with the new Primary. Fast failback works by keeping track of the incoming writes to the new Primary and the writes to the original Primary that did not reach the Secondary before the original Primary failed. Based on the tracked writes, data can be transferred to the original Primary after it recovers. This eliminates the need to completely resynchronize the original Primary with the new Primary after the original Primary recovers; only the changed blocks are resynchronized. Fast failback uses DCMs on the new Primary to track the changed blocks.

To enable fast failback, each data volume on the Secondary must have an associated DCM. The `takeover` command enables fast failback on the new Primary by activating the DCM. The DCM is later used to synchronize the data volumes on the original Primary with the data volumes on the new Primary.

When the original Primary recovers, it needs to be synchronized with the new Primary by playing back the DCM on the new Primary. To receive the missing updates, the original Primary must first be converted to a Secondary. The new Primary can then begin DCM playback to the original Primary. This process can be initiated automatically upon recovery of the original Primary by using the `-autofb` option to the `takeover` command, or it can be initiated manually at some later time by using the `vradmin fbsync` command.

When you use the `-autofb` option with the `vradmin takeover` command, it automatically synchronizes the original Primary when it becomes available. The `-autofb` option converts the original Primary to a Secondary after it comes up and also uses the DCM to synchronize the data volumes on the original Primary using fast failback. The `-autofb` option can be used only if each Secondary data volume has an associated DCM.

If you prefer not to have the original Primary automatically converted to a secondary upon reboot, the process can be performed manually using the `vradmin fbsync` command.

To change the role from Secondary to Primary without enabling fast failback, use the `-N` option with the `vradmin takeover` command. Use the `-N` option if you are sure that the original Primary will not recover or if most of the data on the Primary is going to change while the Primary is down. When performing `vradmin takeover` with the `-N` option

the command automatically detaches the RLINKs from the old Primary to the new Primary. This requires either difference-based synchronization (`vradmin syncrvg`) or full synchronization of the data volumes on the original Primary.

**Example:**

To take over from the Primary RVG `hr_rvg` on host `seattle` to the Secondary RVG `hr_rvg` on host `london`, make sure that the data volumes on the Secondary have associated DCMs. Use the following command to check that the `LOGONLY` attribute is set for the data volumes:

```
# vxprint -g hrdg -ht hr_rvg
```

We recommend that you use the fast failback synchronization method to synchronize the original Primary with the new Primary. For details, see [“Failing Back Using Fast Failback Synchronization”](#) on page 232.

For an RDS with multiple Secondaries, after take over, VVR automatically reconfigures any additional Secondaries as Secondaries of the new Primary, if RLINKs had been created between the Secondaries. Otherwise, you must manually reconfigure them. For more information, see [“Example 2—Taking Over from an Original Primary in a Setup with Multiple Secondaries”](#) on page 229.

To take over an original Primary with fast failback enabled (default), issue the following command on the Secondary that is to take over the Primary role:

```
# vradmin -g diskgroup takeover local_rvgname
```

The argument `diskgroup` is the disk group on the local host.

The argument `local_rvgname` is the name of the RVG on the local host.



## Example 1—Taking Over from an Original Primary

In this example the Primary host `seattle` has failed. This example explains how to take over from the original Primary host `seattle` to the Secondary host `london`. The Secondary host `london` is converted to the new Primary. The disk group name is `hrdg`.

To take over from the Primary `seattle` to Secondary RVG `hr_rvg` on host `london`:

1. Make sure that the data volumes on the Secondary have associated DCMs.

```
# vxprint -g hrdg -ht hr_rvg
```

2. Make the Secondary RVG `hr_rvg` the new Primary RVG by typing the following command on the Secondary `london`:

```
# vradmin -g hrdg takeover hr_rvg
```

The `vradmin takeover` command enables fast failback.

3. Verify whether fast failback is enabled by typing the following command on the Secondary `london`:

```
# vxprint -l rlink_name
```

If fast failback is enabled, the `dcm_logging` flag is set.

4. Start the application on the new Primary `london`. Starting the applications on the new Primary after a takeover may require an application recovery to be run.



## Example 2—Taking Over from an Original Primary in a Setup with Multiple Secondaries

We recommend that you create RLINKs between the hosts `london` and `tokyo` when setting up the RDS.

In this example the Primary host `seattle` has failed. The example explains how to take over from the original Primary host `seattle` to the Secondary host `london`. This example also explains how to start replication from the new Primary `london` to the additional Secondary `tokyo`.

### ▼ To take over from the Primary `seattle` to Secondary RVG on host `london`

1. Make sure that the data volumes on the Secondary have associated DCMs.

```
# vxprint -g hrdg -ht hr_rvg
```

2. Make the Secondary RVG `hr_rvg` the new Primary RVG by typing the following command on the Secondary `london`:

```
# vradmin -g hrdg takeover hr_rvg
```

The `vradmin takeover` command enables fast failback.

3. Verify whether fast failback is enabled by typing the following command on the Secondary `london`:

```
# vxprint -l rlink_name
```

If fast failback is enabled, the `dcm_logging` flag is set.

4. If you had created RLINKs between the Secondary `london` and the additional Secondary `tokyo`, host `tokyo` is automatically added to the new configuration.

Otherwise, you must manually add `tokyo` as a Secondary to the new Primary `london`. To do this, create RLINKs between `london` and `tokyo` and associate them to the respective RVGs using the following commands.

On host `london`:

```
# vxmake -g hrdg rlink rlk_tokyo_hr_rvg local_host=london \
    remote_host=tokyo remote_rlink=rlk_london_hr_rvg \
    remote_dg=hrdg
# vxrlink -g hrdg assoc hr_rvg rlk_tokyo_hr_rvg
```



On host `tokyo`:

```
# vxmake -g hrdg rlink rlk_london_hr_rvg local_host=tokyo \  
    remote_host=london remote_rlink=rlk_tokyo_hr_rvg \  
    remote_dg=hrdg  
  
# vxrlink -g hrdg assoc hr_rvg rlk_london_hr_rvg
```

---

**Note** By default, the `vxmake rlink` command creates the RLINK with the protocol set to UDP/IP. To change the protocol to TCP/IP, see “[Setting the Network Transport Protocol](#)” on page 54.

---

5. Even after takeover, the RLINK from `tokyo` to the original primary `seattle` still remains attached. Detach this RLINK using the following command on the new Primary `london` or on the Secondary `tokyo`:

```
# vradmin -g hrdg stoprep hr_rvg tokyo
```

6. On the new Primary `london`:

- ◆ Synchronize the data volumes in the Secondary RVG `hr_rvg` on `tokyo` with the data volumes in the original Primary RVG `hr_rvg` using the difference-based synchronization and checkpoint. To do this, use the following command on any host in the RDS:

```
# vradmin -g hrdg -c checkpt syncrvg hr_rvg tokyo
```

The `-c` option when used with the `vradmin syncrvg` command automatically starts a checkpoint with the specified name, `checkpt`, in this example. After the data volumes are synchronized the checkpoint is ended.

After the data volumes are synchronized, the checkpoint is ended.

- ◆ Start replication to `tokyo` using the checkpoint created above:

```
# vradmin -g hrdg -c checkpt startrep hr_rvg tokyo
```

7. Start the application on the new Primary `london`. 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 start up to find that one of its Secondaries has been promoted to a Primary by a takeover. This happens when a Secondary of this Primary has taken over the Primary role because of the unexpected outage on this Primary. The process of transferring the role of the Primary back to this original Primary is called *failback*.

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

- ◆ [Failing Back Using Fast Failback Synchronization](#)
- ◆ [Failing Back Using Difference-Based Synchronization](#)

---

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

---

### Fast Failback Versus Difference-Based Synchronization

In the case of fast failback, the data blocks that changed while the original Primary was unavailable are tracked using the DCM for each volume. Difference-based synchronization computes MD5 checksums for a fixed size data block on the Primary and Secondary data volumes, compares it, and then determines whether this data block needs to be transferred from the Primary data volume to the Secondary data volume. The fast failback feature is recommended over the difference-based synchronization for the following reasons:

- ◆ For difference-based synchronization, all the blocks on all the Primary and Secondary data volumes are read; in the case of fast failback, only the blocks that changed on the new Primary are read and hence the number of read operations required is smaller.
- ◆ For difference-based synchronization, the differences are determined by computing and comparing checksum of each of the data chunks on the Secondary and Primary; in the case of fast failback, there is no need to compute checksum because the differences are tracked as they happen, which makes fast failback faster.

The following sections describe each of the above methods for failing back to the original Primary.



## Failing Back Using Fast Failback Synchronization

We recommend that you use the fast failback synchronization method. This procedure assumes that the fast failback feature was enabled on the new Primary when takeover was performed. Failing back to the original Primary using fast failback involves the following steps:

1. Converting the original Primary to an acting Secondary, as shown in the `vxprint -l rvgname` output, and replaying the DCM or SRL of the original Primary to set bits in the DCM of the new Primary. This is performed automatically when the Primary recovers, unless fast failback was disabled during the takeover.

It is possible that the Primary and Secondary data volumes are not up-to-date because all updates to the original Primary might not have reached the Secondary before the takeover. The failback process takes care of these writes by replaying the SRL or DCM of the original Primary. After the original Primary detects that a takeover has occurred, the new Primary uses information in the DCM or SRL of the original Primary to set bits in its DCM for any blocks that changed on the original Primary before the takeover. You can use the `vxlink status` command to monitor the progress of the DCM replay.

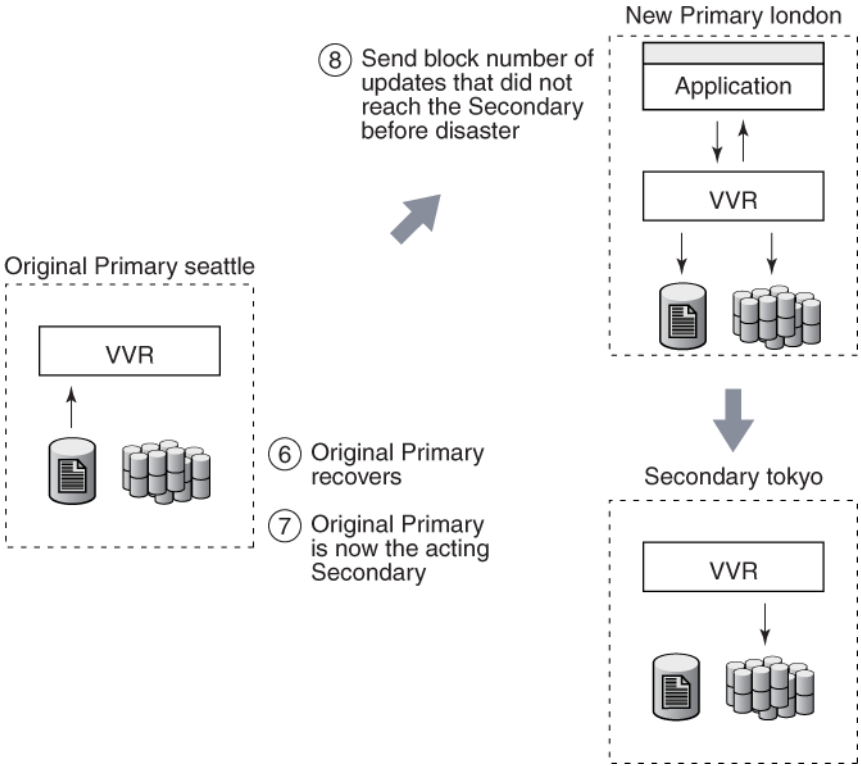
2. Converting the original Primary to Secondary and synchronizing the data volumes on the original Primary with the data volumes on the new Primary using the `vradmin fbsync` command. This command replays the failback log to synchronize the data volumes. The blocks that changed on the original Primary are resynchronized with the new Primary after the DCM of the new Primary is replayed. During the resynchronization, the data from the data volumes on the new Primary is transferred to the data volumes on the original Primary.

This step is not required if the `-autofb` option was used at the time of the takeover. The data on the original Primary data volumes is inconsistent for the duration of the replay. To keep a consistent copy of the original Primary data, take a snapshot of the data volumes before starting the replay. When using the `vradmin fbsync` command, you can also specify the `cache` or the `cachesize` option so that a space-optimized snapshot of the original Primary data volumes is automatically created. If the RVG on the original Primary has VxVM ISP volumes, then you cannot use the `cachesize` attribute.

3. Migrating the Primary Role back to the original Primary and starting replication.

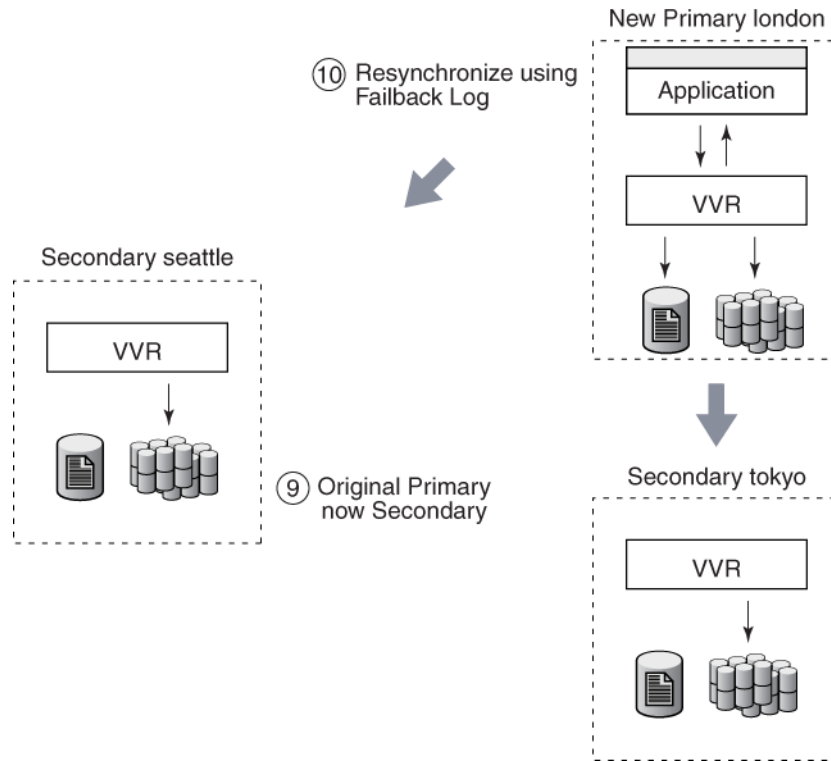


In the following illustration, the original Primary `seattle` has recovered and is now the acting Secondary. The new Primary `london` uses information in the DCM or SRL of the original Primary to set bits in its DCM for any blocks that changed on the original Primary before the takeover.



In the following illustration, the fast failback feature was enabled on the new Primary london when takeover was performed.

The original Primary `seattle` is being resynchronized using the failback log.



### Example 1—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 original Primary `seattle` was taken over by the Secondary host `london`. Each data volume on the Secondary `london` has a Data Change Map (DCM) associated with it. As a result, fast failback is enabled on `london`.

An application is running on `london` and incoming writes are being 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 *seattle* using fast failback

1. Examine the original Primary and make sure you want to convert the original Primary to Secondary.
2. Convert the original Primary to Secondary and synchronize the data volumes in the original Primary RVG *hr\_rvg* with the data volumes on the new Primary RVG *hr\_rvg* on *london* using the fast failback feature. To synchronize the Secondary using fast failback, type the following command on the new Primary *london* or the original Primary *seattle*:

```
# vradm -g hrdg [-wait] fbsync hr_rvg \  
    [cache=cacheobj | cachesize=size]
```

When the synchronization completes, go to the next step. You can check the status of the synchronization using the `vxlink status` command. The `-wait` option with the `vradm fbsync` command can also be used to wait for the completion of the synchronization process.

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. For more information on creating the cache object refer to the information provided in the section [“Preparing the RVG Volumes for Snapshot Operation”](#) on page 164. 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 `vradm fbsync` to create one cache object for each snapshot.

The parameters `cache` and `cachesize` are optional. If you do not specify either of these parameters then the `vradm fbsync` will convert the original Primary to a Secondary and synchronize the data volumes on the original Primary with the data volumes on the new Primary, without creating the snapshots.

This step is not required if the `-autofb` option was used at the time of takeover

3. At a convenient time, stop the application on the new Primary.
4. Migrate the Primary role from the new Primary host *london* to the original Primary host *seattle* by typing the following command on any host in the RDS:

```
# vradm -g hrdg migrate hr_rvg seattle
```

Replication from the original Primary *seattle* to the original Secondary *london* is started by default.

5. Restart the application on the original Primary *seattle*. Because the application was stopped properly before the migration, an application recovery is not required.



## Example 2—Failing Back to the Original Primary Using Fast Failback in a Multiple Secondaries Setup

In this example, the setup consists of two Secondaries, `london` and `tokyo`. The Primary host `seattle` has restarted after an unexpected failure. After the failure, the original Primary `seattle` was taken over by the Secondary host `london`. Each data volume on the Secondary `london` has a Data Change Map (DCM) associated with it. As a result, fast failback is enabled on `london`.

If you had created RLINKs between the hosts `london` and `tokyo` when setting up the replication, you do not need to manually reconfigure the additional Secondary `tokyo` as a part of the RDS. It will automatically be added as a Secondary of the new Primary `london`.

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

1. Perform [step 1 to step 4](#) described in the example [“Example 1—Failing Back to the Original Primary Using Fast Failback”](#) on page 234.
2. After migration, you must synchronize the additional secondary `tokyo` with the original Primary `seattle`.

On the original Primary `seattle`:

- ◆ Synchronize the data volumes in the Secondary RVG `hr_rvg` on `tokyo` with the data volumes in the original Primary RVG `hr_rvg` using the difference-based synchronization and checkpoint. To do this, use the following command on any host in the RDS:

```
# vradmin -g hrdg -c checkpt syncrvg hr_rvg tokyo
```

The `-c` option when used with the `vradmin syncrvg` command automatically starts a checkpoint with the specified name, `checkpt`, in this example. After the data volumes are synchronized the checkpoint is ended.

- ◆ Start replication to `tokyo` using the checkpoint created above:

```
# vradmin -g hrdg -c checkpt startrep hr_rvg tokyo
```

3. Restart the application on the original Primary `seattle`. Because the application was stopped properly before the migration, an application recovery is not required.

## Failing Back Using Difference-Based Synchronization

Failing back to the original Primary using difference-based synchronization involves the following steps:

1. Converting the original Primary to a Secondary of the new Primary.
2. Synchronizing the data volumes on the original Primary with the data volumes on the new Primary using difference-based synchronization using checkpoint.
3. Starting replication to the Secondary (original Primary) using checkpoint.
4. Migrating the Primary Role to the original Primary and starting replication by default.

The examples in the following section explain how to fail back to the original Primary using VVR.

### Converting an Original Primary to a Secondary

VVR provides the `vradmin makesec` command to convert the original Primary to a Secondary. This command is only needed if fast failback was not enabled when the original takeover was performed. If fast failback was enabled, the original Primary will automatically be converted to a Secondary when DCM playback begins. Note that this command can be run only from the original Primary host where one of its original Secondaries has taken over the Primary role.

You can issue the `vradmin makesec` command in the failback procedure only if fast failback has not been enabled when taking over from the original Primary. Run this command when the original Primary restarts. Stop the application if it restarts automatically when the Primary restarts. The `vradmin makesec` command converts the original Primary to a Secondary RVG.

---

**Tip** Before using the `vradmin makesec` command make sure you close all the applications running on the original Primary's data volumes. Also ensure that none of the data volumes are open.

---

The `vradmin makesec` command fails if the Secondary data volumes are not up-to-date or if there are some applications still running on the failed Primary's data volumes. Use the `-f` option to convert a failed Primary to a Secondary even when the Secondary data volumes are not up-to-date. If any of the failed Primary data volumes are open or have applications running on them, then using the `vradmin makesec` command with the `-f` option will fail. To proceed with the `vradmin makesec` command, first close the volumes or stop the applications as required.



To convert an original Primary to a Secondary:

```
# vradmin -g diskgroup makesec local_rvgname newprimary_name
```

The argument *diskgroup* is the disk group on the local host.

The argument *local\_rvgname* is the name of the RVG on the local host, that is, the original Primary and represents its RDS.

The argument *newprimary\_name* is the name of the new Primary host, that is, the previous Secondary host. Note that the *newprimary\_name* argument must be the same as the host name displayed with the Primary-Primary configuration error in the output of the `vradmin -l printrvg` command.

### Example 3—Failing Back to the Original Primary Using Difference-Based Synchronization

In this example, the Primary host `seattle` has restarted after an unexpected failure. After the failure, the original Primary `seattle` has been manually taken over by the Secondary host `london`. This example shows how to fail back to the original Primary `seattle` using difference-based synchronization. For more information, see [“Synchronizing Volumes Using Difference-Based Synchronization”](#) on page 125.

#### ▼ To fail back to the original Primary `seattle`

1. Make the original Primary RVG `hr_rvg` on `seattle` the Secondary RVG of the new Primary `london` by typing the following command on the original Primary `seattle`:

```
# vradmin -g hrdg makesec hr_rvg london
```

2. Synchronize the data volumes in the original Primary RVG `hr_rvg` with the data volumes in the new Primary RVG `hr_rvg` on `london` using the difference-based synchronization and checkpoint. To synchronize the Secondary based on differences using a checkpoint, type the following command on any host in the RDS:

```
# vradmin -g hrdg -c checkpoint_presync syncrvg hr_rvg seattle
```

3. Stop the application on the new Primary `london`.
4. Start replication to the Secondary RVG (original Primary) `hr_rvg` on `seattle` from the new Primary RVG `hr_rvg` on `london` using the checkpoint by typing the following command on any host in the RDS:

```
# vradmin -g hrdg -c checkpoint_presync startrep hr_rvg seattle
```

5. Migrate the Primary role from the new Primary host `london` to the original Primary host `seattle` by typing the following command on any host in the RDS:
 

```
# vradmin -g hrdg migrate hr_rvg seattle
```

 Replication from the original Primary `seattle` to the original Secondary `london` is started by default.
6. Restart the application on the original Primary `seattle`. Because the application was stopped properly before the migration, an application recovery is not required.

### Example 4—Failing Back to the Original Primary Using Difference-Based Synchronization in a Multiple Secondaries Setup

This example shows how to fail back to the original Primary `seattle` using the difference-based synchronization feature.

1. Perform [step 1](#) to [step 5](#) described in the example “[Example 3—Failing Back to the Original Primary Using Difference-Based Synchronization](#)” on page 238.
2. On the original Primary `seattle`:
  - ◆ Synchronize the data volumes in the Secondary RVG `hr_rvg` on `tokyo` with the data volumes in the original Primary RVG `hr_rvg` using the difference-based synchronization and checkpoint. To do this, use the following command on any host in the RDS:
 

```
# vradmin -g hrdg -c checkpt syncrvg hr_rvg tokyo
```

 The `-c` option when used with the `vradmin syncrvg` command automatically starts a checkpoint with the specified name, `checkpt`, in this example. After the data volumes are synchronized the checkpoint is ended.
 

After the data volumes are synchronized, the checkpoint is ended.
  - ◆ Start replication from `seattle` to `tokyo` using the command:
 

```
# vradmin -g hrdg -c checkpt startrep hr_rvg tokyo
```
3. Restart the application on the original Primary `seattle`. Because the application was stopped properly before the migration, an application recovery is not required.







# Administering VVR Using VVR VEA

---

# 8

VERITAS Enterprise Administrator (VEA) is a Java-based graphical user interface that can be used to configure and manage storage objects. VEA enables you to access multiple hosts and perform various product-specific tasks. VEA consists of a VEA server and a VEA client. The VEA client communicates with the server to display objects of the products residing on the server and perform operations on these objects. The server runs on a system that runs VERITAS Volume Replicator (VVR). The client runs on any platform that supports the Java Runtime Environment.

VVR VEA enables you to configure, monitor, and administer VVR in a distributed environment. Using VVR VEA, you can manage VVR objects on multiple hosts. VVR VEA also enables you to perform VVR tasks from any host in the VVR environment.

For detailed information on how to get started and work with VERITAS Enterprise Administrator, refer to *VERITAS Enterprise Administrator (VEA 500 Series) Getting Started*.

Refer to the VxVM VEA online help for information on how to use VEA to perform VxVM tasks. For detailed information about VVR, refer to the appropriate chapters in the *VERITAS Volume Replicator Administrator's Guide* (this guide).



## Features of VVR VEA

VVR VEA provides the following features:

- ◆ Remote Administration

You can administer VVR remotely or locally using the VEA client, which runs on UNIX or Windows machines. This GUI enables you to administer multiple hosts simultaneously.

- ◆ Wizards and Menus

VVR VEA offers wizards and menus that guide you through complex VVR administrative tasks, including creating an RDS, starting replication, and transferring the Primary role.

- ◆ Detailed Views

VVR VEA provides detailed views that display information about the selected VVR objects. For example:

The RDS view displays detailed information about the selected Replicated Data Set (RDS).

The Alerts view displays alerts for important VVR events along with a detailed description of the event, its causes, and recommended actions.

- ◆ Mechanism for Notifying Users

You can receive SNMP notifications or Email notifications of any alarms or events related to VVR. You can also specify the commands or scripts to be executed automatically on detecting an event.

## Replicated Data Set Defined

Volume Replicator (VVR) replicates data from a Primary host, where the application is running, to one or more Secondary hosts. A Replicated Volume Group (RVG) on the Primary host and its counterparts on the Secondary hosts make up a Replicated Data Set (RDS). An RDS is not a VERITAS Volume Manager object; it is a concept used in VVR. An RDS enables grouping of the RVG on the Primary host and its counterparts on the Secondary host.

VVR VEA operates on an RDS, that is, the Primary RVG and all the Secondary RVGs in the RDS. Using VVR VEA, you can perform tasks from any host in an RDS; VVR performs the same tasks on the required hosts in the RDS if appropriate.



## Getting Started with VVR VEA

This section explains how to start VVR VEA from a Windows client and from a UNIX client.

### Launching VEA

You can start the VEA console from a UNIX or Windows client.

#### From a UNIX Client

To start VEA from the command line:

1. Go to the directory `/opt/VRTSob/bin` where the VEA client is installed.
2. Start the VEA client by running the `vea` script from this location. While starting the VEA client from the command line, the following options are available:
  - v: Shows the version of client console
  - c: Specifies the configuration file for advanced settings
  - host: Specifies the host to connect to. If the username and password for the host is already stored, the stored information is used by VEA; otherwise, VEA prompts you to enter the user name and password.

#### From a Windows Client

From the Windows client, you can start VEA either from the Start menu or from the command line.

##### ▼ To start VEA from the Start menu

Choose **Start > Programs > VERITAS > VERITAS Enterprise Administrator**.

##### ▼ To start VEA from the command line

1. Go to the directory where the VEA client is installed. For example,  
`c:\Program Files\VERITAS\VERITAS Object Bus\Bin`
2. Start the VEA client by running the `vea.exe` file from this location. While starting the VEA client from the command line, the following options are available:
  - v: Shows the version of client console



-c: Specifies the configuration file for advanced settings

-host: Specifies the host to connect to. If the username and password for the host is already stored, the stored information is used by VEA; otherwise, VEA prompts you to enter the user name and password.

## Connecting to the Hosts

Choose **File > Connect** and complete the Connection dialog box to connect to the host on which you want to create or view VVR configurations.

## Closing VVR

To exit the VVR VEA, do the following:

Choose **File > Exit** or click the close box in the upper right corner of the VEA console.

## Working with VVR

VVR VEA includes the following components:

- ◆ [VVR VEA Main Window](#)
- ◆ [Tree and Detailed Views](#)
- ◆ [Replication Menu on the Menu Bar](#)
- ◆ [Pop-up Menus](#)

Note that explanation of each component is specific to VVR. For more information about these components, see the *VERITAS Volume Manager User's Guide-VEA*.



## VVR VEA Main Window

To use the VVR features in VERITAS Enterprise Administrator, enable the optional VVR functionality on the hosts. To enable the VVR feature, install a VVR license on each host on which the VEA server runs.

When you connect to a server that has the VVR license enabled, the node Replicated Data Sets appears under the host name in the tree. The main window consists of the following panes:

Left pane—Displays the tree view of all objects that are managed using VEA.

Right pane—Displays views of the objects selected in the tree.

Lower pane—Consists of a Console View and Tasks View

## Tree and Detailed Views

### Left Pane (Tree)

The left pane contains the *tree*, which is a dynamic hierarchical display of the VERITAS Volume Manager objects, VVR objects and components, and other objects on the selected host. Each node in the tree represents a group of objects. Each object group consists of an icon and a name.

In the tree, the **Replicated Data Sets** node displays the VVR components. The Replicated Data Sets node displays all the RDSs on the host.

- ◆ To display all the RDSs on the selected host, expand the Replicated Data Sets node.
- ◆ To hide the RDSs, collapse the Replicated Data Sets node.

The following nodes appear when you expand the RDS name node:

- ◆ Primary host name
- ◆ Secondary host name

The name of the Primary host appears below the name of the RDS followed by the name of the Secondary host. If an RDS contains multiple Secondary hosts, the names of all the Secondary hosts are displayed, one below the other, at the same level.

---

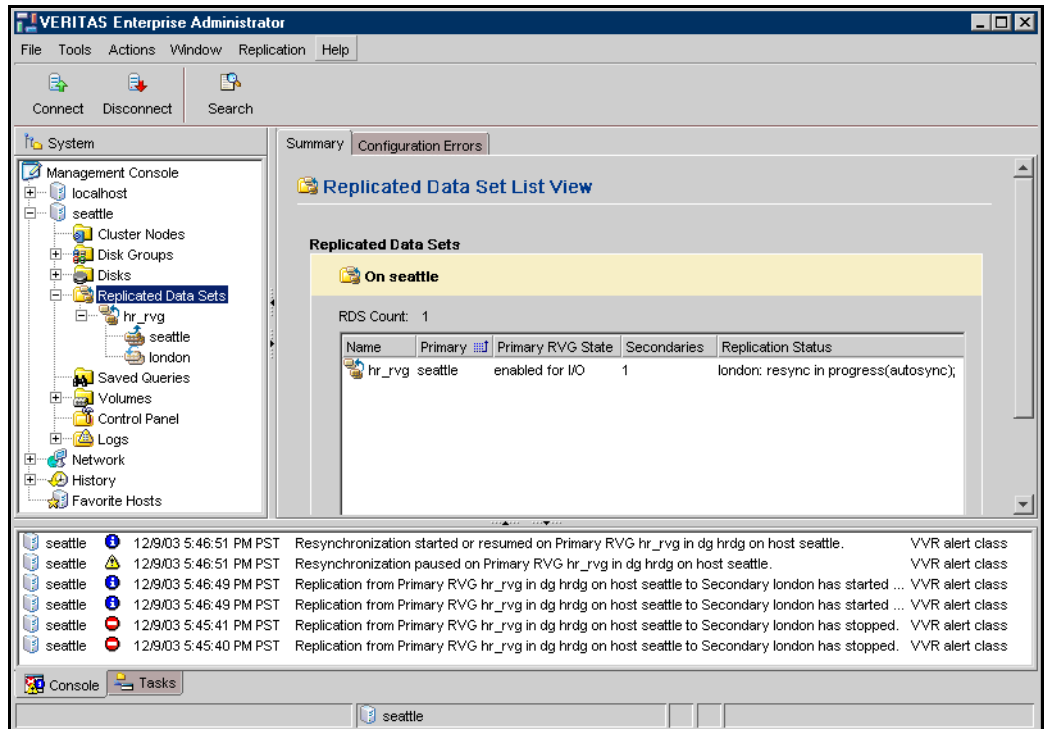
**Note** The Replicated Data Sets node appears in the tree view only if you have a VVR license. When you use VVR, use the Replicated Data Sets node to perform VVR tasks. Use the VM tasks to create and view volumes.

---

## Right Pane (Detailed Views)

The right pane displays detailed views of the object that is selected in the tree view or the navigational view. The right pane could consist of multiple tabbed views that display detailed configuration information and status information.

### Right Pane (Detailed View)



The detailed view also contains a *grid*, which is a tabular display of the properties of the selected object. You can perform operations on the objects displayed in the grid. For example, you can associate a volume to an RDS by selecting the name of the RDS in the grid, and then choosing **Replication > Associate Volume**. The grid constantly updates its contents to reflect changes to objects.

The *splitter* is the vertical bar that separates the tree from the configuration view (grid). To resize the left pane and the right pane, hold the mouse button over the splitter and drag it to the left or right.



### Lower Pane (Status Views)

The lower pane provides two views that are tabbed.

**Console View**—Displays events and alerts for the connected hosts. Double-clicking an alert displays more information about the alert.

**Tasks View**—Displays information related to tasks that were run. For example, the Tasks View displays the time the task was started, object name, and so on.

---

**Note** When a client Java exception or error occurs, an error alert icon is displayed in the lower region of the main window. Click the icon to display the error or exceptions. For more information on troubleshooting client problems, refer to the *VERITAS Volume Manager User's Guide—VEA*.

---

## Replication Menu on the Menu Bar

The Replication menu on the menu bar is a context-sensitive menu that launches tasks on a selected object. The Replication menu contains a submenu that enables you to perform VVR tasks. The Replication menu is dynamic and changes its options based on the object that is selected. By default, the Replication menu is hidden.

When a VVR object is selected in either the left or the right pane, the Replication menu is updated with submenus that enable you to perform tasks that are appropriate for the selected object. For example, the Replication menu shows the Create Replicated Data Set submenu when you select the Replicated Data Sets node in the tree view.

## Pop-up Menus

In addition to the Replication Menu in the Main Menu, VVR VEA provides pop-up menus that enable you to select tasks and properties. To select a task from the pop-up menu, do one of the following:

- ❖ Select the node in the left pane or tree view, for example, Replicated Data Sets or a Replicated Data Set, right-click to display the pop-up menu, and then select the required task.
- ❖ Select the icon in the left pane or tree view, for example, Replicated Data Sets, select an object from the grid, right-click to display the pop-up menu, and then select a task.



## Setting Up Replication Using VVR VEA

VVR VEA enables you to set up Volume Replicator (VVR) configurations on participating hosts from a remote client. VVR enables you to set up replication either when the data volumes are zero initialized or contain valid data, and when the application is active or stopped.

Before setting up a Replicated Data Set, lay out your VVR configuration using the Configuration Worksheet in the *Volume Replicator Installation Guide*. Follow the best practices or recommendations listed in “[Best Practices for Setting Up Replication](#)” on page 42 to ensure successful configuration of VVR. You can set up replication using one of the following procedures:

- ✓ [Setting Up Replication Using the Create RDS Wizard](#)
- ✓ [Setting Up Replication Using Individual Tasks](#)

The **Create Replicated Data Set** wizard guides you through the process of setting up replication using automatic synchronization. For instructions on setting up replication using the **Create Replicated Data Set** wizard, see “[Setting Up Replication Using the Create RDS Wizard](#)” on page 252. For instructions on setting up replication using separate tasks, see “[Setting Up Replication Using Individual Tasks](#)” on page 258.

## Overview of Setting Up Replication

To configure and set up replication, the following tasks must be performed in the order presented below:

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

---

VVR VEA enables you to create a Primary RVG of an RDS using the **Create Replicated Data Set** wizard. Launch the wizard from the host that you want to be designated as the Primary in the RDS.



The **Create Replicated Data Set** wizard enables you to associate existing data volumes and the Storage Replicator Log (SRL) with the Primary RVG. The **Create Replicated Data Set** wizard performs the following operations:

- ◆ Creates the Primary RVG on the host to which you connected.
- ◆ Enables the Primary RVG by starting it.
- ◆ Associates the specified data volumes and Storage Replicator Log (SRL) with the RVG.
- ◆ Associates a Data Change Map (DCM) to each data volume in the RVG.

You can choose to add the Secondary using this wizard or add the Secondary later using the Add Secondary wizard.

You can choose to synchronize the Secondary later using the Synchronize Secondary wizard. Alternatively, you can synchronize the Secondary using the Start Replication dialog box and by choosing either Automatic Synchronization, Synchronization not needed, or Checkpoint options.

If an operation fails, a failure message lists the possible reasons of failure in the **Result Details** area. Read the details carefully and take appropriate corrective action. For more information about the error, note down the Unique Message Identifier (UMI) that is displayed with the error message, and then refer the VERITAS Technical Services knowledge base.

---

**Note** UMI is a string of the format "V-xx-yy-zz". For example, V-38-1-11 is the UMI displayed when an invalid Secondary host name is entered.

---

## Example Scenario

The instructions on how to use VVR VEA to set up a simple VVR configuration under different situations include examples. These examples explain how to set up a VVR configuration containing one Primary and one Secondary. Note that VVR VEA enables you to create and set up VVR configurations containing multiple Secondary hosts. The example configuration looks as follows:

Primary Host name: `seattle`

Disk group	<code>hrdg</code>
Primary RVG	<code>hr_rvg</code>
Primary RLINK to Secondary <code>seattle</code>	<code>rlk_seattle_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>london</code>	<code>rlk_london_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>

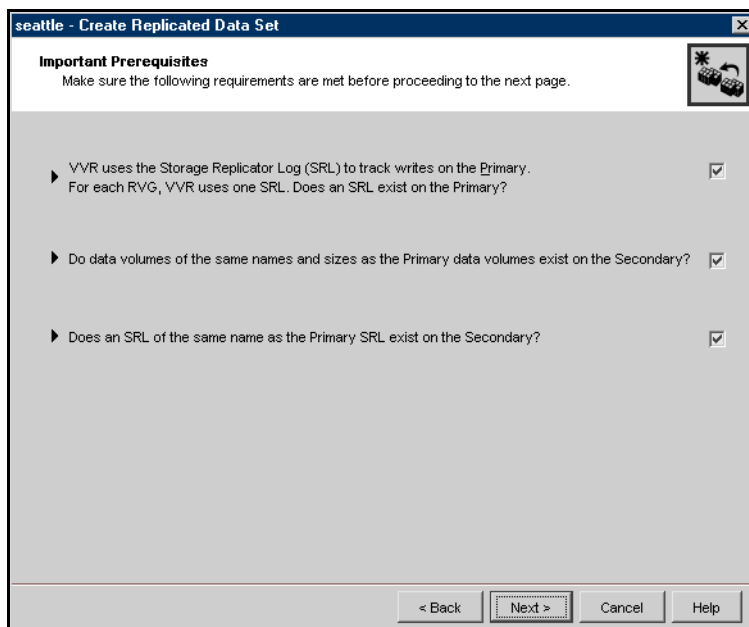


## Setting Up Replication Using the Create RDS Wizard

The **Create Replicated Data Set** wizard guides you through the process of setting up replication using automatic synchronization.

### ▼ To set up replication using the Create Replicated Data Set wizard

1. Launch the wizard from the host that you want to be designated as the Primary in the RDS.
2. Choose **Replication > Create Replicated Data Set**. To use the pop-up menu, right-click **Replicated Data Sets** and select **Create Replicated Data Set Wizard**.
3. In the Replicated Data Set Creation Wizard page, select **Beginner User** or **Advanced User**.
4. If you selected **Advanced User**, skip to step 6. If you selected **Beginner User**, the Important Prerequisites page is displayed.
5. In the Important Prerequisites page, click the check boxes and then click **Next**. Note that this page does not appear if you selected **Advanced User** in step 2. You cannot proceed if the data volumes and SRL of the same names and sizes do not exist on all the hosts in the RDS.



6. Complete the Assign a name to the Primary RVG page as follows:

Primary RVG Name	Enter the name of the Primary RVG. The name of the Primary RVG is designated as the name of the RDS. For example, <code>hr_rvg</code> .
Include VxVM ISP Volumes Only	Select this checkbox only if you want to include VxVM ISP (Intelligent Storage Provisioning) volumes in the RVG. You <i>cannot</i> include both VxVM ISP volumes and VxVM non-ISP volumes in the RVG.

After you have provided the required information, click **Next**.

7. Complete the Select the Disk Group and Data Volumes to be replicated page as follows:

Disk Group	Select the required disk group name from the drop-down list. For example, <code>hr_dg</code> .
Data Volumes	Select the names of the data volumes by holding down the SHIFT key and clicking the names of the volumes you want to select. For example, <code>hr_dv01</code> , <code>hr_dv02</code> . Note that volumes with the same names must be created on the Primary host and all Secondary hosts. If you selected <b>Include VxVM ISP Volumes Only</b> , the list displays VxVM ISP volumes only.

After you have provided the required information, click **Next**.

8. In the Select the Replication Log page, select the name of the SRL, for example, **hr\_srl**, and then click **Next**. If you selected **Include VxVM ISP Volumes Only**, the list displays VxVM ISP volumes only.
9. In the Summary Information for Creating the Primary Replicated Volume Group page, select Create Primary RVG after reviewing the information. If the operation of creating Primary RVG was successful, the next page displays a message that indicates that the Primary RVG has been created successfully.

If the operation fails, a failure message lists the possible reasons of failure in the **Result Details** area. Read the details carefully and take appropriate corrective action. For more information about the error, note down the Unique Message Identifier (UMI) that is displayed with the error message, and then refer the VERITAS Technical Services knowledge base.



**Note** UMI is a string of the format "V-xx-yy-zz". For example, V-38-1-11 is the UMI displayed when an invalid Secondary host name is entered.

---

**10.** To add the Secondary now, click the **Do you want to add Secondary?** check box. If you want to add the Secondary later, click **Finish**.

**11.** Complete the Add Secondary Wizard page as follows:

<b>Hostname/IP Address for Replication</b>	
Primary	Enter a host name or specify the IP address that can be used for replication. The Primary host name must be resolvable and reachable from the Secondary host. For example, <i>seattle</i> . If you entered the IP address, it must be reachable from the Secondary host. The Replication IP column in the Secondary RVG view displays the value of the Primary host.
Secondary	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, <i>london</i> . If you entered the IP address, it must be reachable from the Primary host. The Replication IP column in the Primary RVG view displays the value of the Secondary host.
Replication Setting	The default values for Replication Mode, SRL Protection, and Protocol are displayed. You can keep the default or change these replication settings. To change the replication settings, click <b>Advanced</b> . The Advanced Settings dialog box appears.

12. To change the replication settings, complete the Advanced Settings dialog box as follows, or go to the next step.

**Advanced Settings**

The following parameters are optional. If you do not change any of these parameters the default values will be used.

Secondary DG :

Primary RLINK :       Secondary RLINK :

Replication Settings:

Replication Mode:  Asynchronous  
 Synchronous-Override

SRL Protection :  AutoDCM      Latency Protection :  Off     Override  
 DCM  
 Override  
 Off

Latency Low Mark:   
Latency High Mark:

Packet Size:       Protocol :  UDP     TCP

Set Bandwidth Limit      Bandwidth Limit:    

OK    Cancel    Help

Secondary DG	If the name of the Secondary disk group is other than the default, enter the correct 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.
Primary RLINK	Enter the new Primary RLINK name if the default name is not appropriate.
Secondary RLINK	Enter the new Secondary RLINK name if the default name is not appropriate.
Replication Settings	



Replication Mode	Asynchronous is selected by default. Select the required mode of replication. To enable synchronous mode, select <b>Synchronous-Override</b> . To set the replication mode to <code>synchronous=fail</code> , use the command-line interface. For more information, refer to “ <a href="#">Administering VERITAS Volume Replicator</a> ” on page 121.
SRL Protection <b>AutoDCM, DCM, Override, Off</b>	By default, SRL protection is set to <b>AutoDCM</b> . Select the required value to change the SRL protection mode. To enable SRL protection, select <b>AutoDCM, DCM, or Override</b> . To disable SRL protection, select <b>Off</b> . If you disable SRL protection, the SRL is allowed to overflow. To set the SRL protection mode to <b>fail</b> , use the command-line interface (CLI). For details on using the CLI, refer to “ <a href="#">Administering VERITAS Volume Replicator</a> ” on page 121.
Latency Protection <b>Override, Off</b>	By default, latency protection is set to <b>Off</b> . To enable latency protection and limit the number of updates by which the Secondary is allowed to fall behind, select <b>Override</b> . If you select <b>Override</b> , 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 <b>Off</b> .  To set the latency protection to <b>fail</b> , use the command-line interface (CLI). For details on using the CLI, refer to “ <a href="#">Administering VERITAS Volume Replicator</a> ” on page 121.
If Latency Protection is set to <b>Override</b> :	
Latency High Mark	Displays only if Latency Protection is set to <b>Override</b> . The default is 10000. Enter the required numerical limit.
Latency Low Mark	Displays only if Latency Protection is set to <b>Override</b> . The default is 9950. Enter the required numerical limit.
Packet Size	The default is 8400. Enter the required packet size.
Protocol	UDP is selected by default. Select UDP or TCP.
Set Bandwidth Limit	To control the network bandwidth to be used by VVR for replication, click the <b>Set Bandwidth Limit</b> check box; then specify the bandwidth limit. Note that the specified bandwidth limit must be greater than 56 kilo bits per second (kbps).  To disable Bandwidth Throttling, clear the <b>Set Bandwidth Limit</b> check box. This sets the bandwidth limit to <b>none</b> . By default, Bandwidth Throttling is disabled.



13. Select **OK**, and then select **Add Secondary**.
14. In the next page, the message indicates that the Secondary has been added successfully. If the operation of adding Secondary RVG was successful the page displays a success message.  
  
If the operation fails, the page displays the failure message and lists the possible reasons of failure in the **Result Details** area. Read the details carefully and take appropriate corrective action. For more information about the error, note down the Unique Message Identifier (UMI) that is displayed with the error message, and then refer the VERITAS Technical Services knowledge base.
15. Select **Do you want to synchronize Secondary?** to synchronize the Secondary using automatic synchronization; otherwise, to quit the wizard, click **Finish**.
16. In the Synchronize Secondary Using Automatic Synchronization page, select **Synchronize**, and then select **Finish**.

**Note:**

- ◆ To add more volumes after creating the RVG, use the **Associate Volume** task.



## Setting Up Replication Using Individual Tasks

To configure and set up replication without completely following through the Create Replicated Data Set wizard, perform the following discrete tasks:

- ✓ [Creating a Replicated Data Set](#)
- ✓ [Synchronizing the Secondary and Starting Replication](#)

### Creating a Replicated Data Set

To create a Replicated Data Set (RDS) using VVR VEA, 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

Prerequisites for creating a Primary RVG:

- ✓ Data volumes of the same names and size must exist on all the hosts that are to be a part of the RDS.
- ✓ SRLs of the same name must exist on all the hosts that are to be a part of the RDS.

#### ▼ To create a Primary RVG of an RDS

1. Perform [step 1 to step 9](#) in “[Setting Up Replication Using the Create RDS Wizard](#)” on page 252.
2. Click **Finish** to create the Primary RVG of the RDS and quit the Create Replicated Data Set wizard.

#### Adding a Secondary to the RDS

#### ▼ To add a Secondary to the RDS

1. From the tree view, select the name of the RDS for which you want to create the Secondary. For example, **hr\_rvg**.
2. Choose **Replication > Add Secondary**. To use the pop-up menu, right-click the name of the RDS, for example, **seattle**. The Add Secondary Wizard is displayed.
3. Click Next.

4. Perform [step 11](#) to [step 14](#) in “[Setting Up Replication Using the Create RDS Wizard](#)” on page 252.
5. Click **Finish** to add the Secondary to the RDS.



## Synchronizing the Secondary and Starting Replication

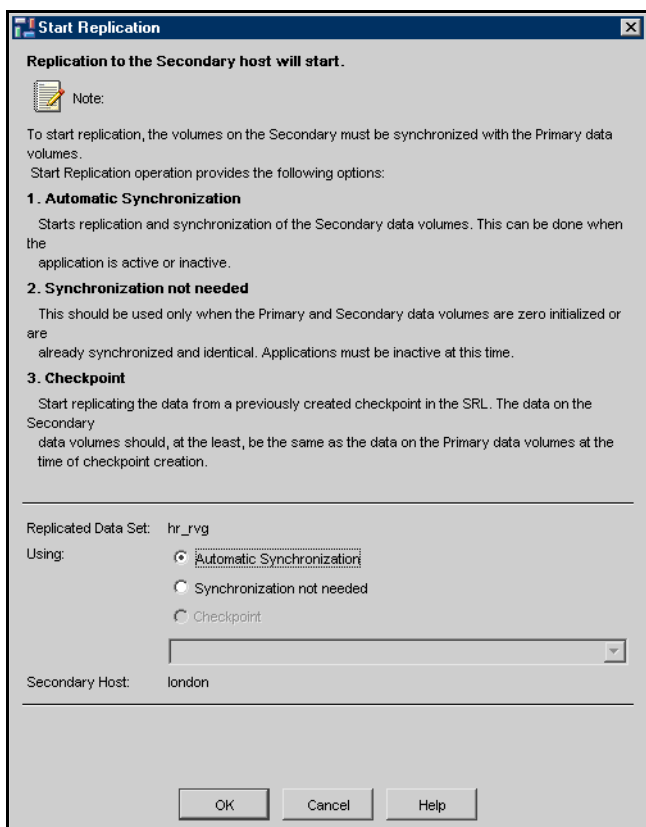
Before starting replication, the Secondary data volumes must be synchronized with the Primary data volumes. The Start Replication task enables you to synchronize the Secondary and start replication between the Primary and a specific Secondary of an RDS.

The Start Replication task can be launched from any host in an RDS. The name of the RVG on the local host represents the RDS.

### To synchronize the Secondary and start replication:

1. Select the name of the Secondary that is to be synchronized and to which you want to start replication. For example, **london**.
2. Choose **Replication > Start Replication**. The [Start Replication Dialog Box](#) is displayed.

Start Replication dialog box



## Start Replication Dialog Box

The Start Replication task provides the following options to start replication:

- ◆ Automatic Synchronization—For instructions, see [“Starting Replication Using Automatic Synchronization”](#) on page 261.
- ◆ Synchronization not needed—For instructions, see [“Starting Replication When Data Volumes are Initialized with Zeroes”](#) on page 264.
- ◆ Checkpoint—For instructions, see [“Setting Up Replication Using Checkpoint”](#) on page 262.

### Notes:

- ◆ Depending on the selected option, the Start Replication task synchronizes the Secondary data volumes with the Primary data volumes, and then starts replication. For more information on the methods to synchronize the Secondary and how to choose the appropriate method to synchronize the Secondary, see [“Methods to Synchronize the Secondary”](#) on page 58.
- ◆ The **Synchronize Secondary** wizard synchronizes the Primary and Secondary data volumes only.

## Starting Replication Using Automatic Synchronization

The instructions in this section assume that the RDS has been created using the procedure described in the section [“Creating a Replicated Data Set”](#) on page 258.

### ▼ To synchronize the Secondary and start replication using automatic synchronization

1. Select the name of the Secondary that is to be synchronized and to which you want to start replication. For example, **london**.
2. Choose **Replication > Start Replication**. To use the pop-up menu, right-click the name of the Secondary, for example, **london**. The Start Replication dialog box is displayed.



3. From the Start Replication dialog box, select **Automatic Synchronization** and click **OK**. The Start Replication dialog box displays the result of your request to start replication.

**Notes:**

- While automatic synchronization is in progress, the Replicated Data Set View for the RDS displays the Data Status as **inconsistent**, and the Replication Status as **resync in progress (autosync)**.
- After the automatic synchronization completes, the Data Status changes to **consistent; up-to-date** and the Replication Status changes to **replicating (connected)** in the Replicated Data Set View for the RDS.

## Setting Up Replication Using Checkpoint

The instructions in this section assume that the RDS has been created using the procedure described in the section “[Creating a Replicated Data Set](#)” on page 258. You can synchronize the Secondary using block-level backup and checkpoint when the application is active or inactive.

1. Start a checkpoint using the command-line interface. For instructions, see “[Creating Checkpoints](#)” on page 154.
2. Perform a block-level backup of the data volumes in the Primary RVG.
3. End the checkpoint using the command-line interface.
4. Restore the backup to the Secondary data volumes.
5. Check whether the checkpoint you created is still valid by selecting the Checkpoints tab in the Primary RVG view and then clicking **Get Checkpoints**. If the checkpoint has overflowed, repeat [step 1](#) to [step 6](#).
6. Choose **Replication > Start Replication** to start replication. The Start Replication dialog box is displayed.
7. From the Start Replication dialog box, select **Checkpoint** to synchronize the Secondary and start replication using checkpoint.
8. From the drop-down box, select the name of the checkpoint.
9. Click **OK**. The Start Replication dialog box displays the result of your request to start replication.



10. Click **OK** to close the Results dialog box. The RDS View for `hr_rvg` on `seattle` is automatically refreshed to display the replication status at a particular point in time, for the replication that has just been started. If the Secondary is consistent, the synchronization was successful.

If the checkpoint overflows before the Secondary becomes consistent, the synchronization fails. Increase the size of the SRL, and then repeat [step 1](#) to [step 11](#). For instructions on resizing the SRL, see “[Resizing the SRL](#)” on page 141.

11. In the RDS view, check Data Status to find out whether the Secondary is up to date.

**Notes:**

- ◆ 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.
- ◆ While synchronization using checkpoint is in progress, the Replicated Data Set View for the RDS displays the Data Status as **inconsistent**, and the Replication Status as **replicating (connected)**.
- ◆ After the synchronization using checkpoint completes, the Data Status changes to **consistent** and the Replication Status changes to **replicating (connected)** in the Replicated Data Set View for the RDS.



## Starting Replication When Data Volumes are Initialized with Zeroes

Because the Primary data volumes are initialized with zeroes, the data on the Secondary node does not need to 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 the section “[Creating a Replicated Data Set](#)” on page 258.

---

**Note** Do not start the application or mount the file system before completing [step 4](#).

---

1. Select the Secondary host name from the tree, and then choose **Replication > Start Replication**.
2. From the Start Replication dialog box, select **Synchronization not needed**.
3. Click **OK**. The Start Replication dialog box displays the result of your request to start replication.
4. Click **OK** to close the Start Replication dialog box.
5. Start the application.

<b>Notes:</b>
<ul style="list-style-type: none"><li>◆ When this task is initiated, the Replicated Data Set View for the RDS displays the Data Status as <b>consistent</b>, and the Replication Status as <b>replicating (connected)</b>.</li><li>◆ When this task completes, the Replicated Data Set View for the RDS displays the Data Status as <b>consistent</b>, and the Replication Status as <b>replicating (connected)</b>.</li></ul>



## Viewing Replication Information

You can get information about the VVR environment using the various views presented by VVR VEA. VVR VEA provides a tree view, which is a hierarchical top-level view for each VVR object. The tree view enables you to drill down to get detailed information for each VVR object. For each VVR object in the tree view, VVR VEA provides an associated view that consists of tabbed views containing related information about the selected node. You can display the VVR VEA views on any host in an RDS.

### Viewing the RDSs Present on the Selected Host

When you click the **Replicated Data Sets** node in the tree view, the RDSs that are configured on the selected hosts are displayed in the tree view; the right pane displays the RDS List View, which consists of the Summary tab and the Configuration Errors tab.

- ✓ The Summary tab displays information about the RDSs present on the selected host, such as the number of RDSs of which the selected host is a part, name of the RDS, name of the Primary, Primary RVG state, number of Secondaries in the RDS, and the replication status of the RDS.
- ✓ To view configuration errors, if any, in the RDS use the Configuration Errors tab. Each configuration error in an RDS is displayed on a separate line. If multiple configuration errors exist in an RDS, the view displays multiple lines for that RDS.
- ✓ The Configuration Errors view is persistent across all the views in VVR VEA and displays configuration errors related to the object selected in the tree view. When the node selected in the tree view is the **Replicated Data Sets** node, the Configuration Errors view displays configuration errors for all the RDSs. Similarly, if the node selected in the tree view is an RDS, the Configuration Errors view displays errors related to that particular RDS only.

### Using the RDS List View

- ◆ To display the RDS List View from any other view, click the **Replicated Data Sets** node in the tree view. You can go to the RDS View from the RDS List View, by double-clicking the RDS name or by clicking the RDS node in the tree view.



## RDS List View (Summary tab)

The screenshot displays the VERITAS Enterprise Administrator interface. The left-hand pane shows a tree view of the system hierarchy, with 'Replicated Data Sets' expanded to show 'hr\_rvg' on 'seattle'. The main pane is titled 'Replicated Data Set List View' and shows a table of replicated data sets. Below the table, a console window displays a series of alerts related to replication synchronization between 'seattle' and 'london'.

**Replicated Data Sets**

**On seattle**

RDS Count: 1

Name	Primary	Primary RVG State	Secondaries	Replication Status
hr_rvg	seattle	enabled for I/O	1	london: resync in progress(autosync);

Console Log:

- seattle 12/9/03 5:46:51 PM PST Resynchronization started or resumed on Primary RVG hr\_rvg in dg hrdg on host seattle. VVR alert class
- seattle 12/9/03 5:46:51 PM PST Resynchronization paused on Primary RVG hr\_rvg in dg hrdg on host seattle. VVR alert class
- seattle 12/9/03 5:46:49 PM PST Replication from Primary RVG hr\_rvg in dg hrdg on host seattle to Secondary london has started ... VVR alert class
- seattle 12/9/03 5:46:49 PM PST Replication from Primary RVG hr\_rvg in dg hrdg on host seattle to Secondary london has started ... VVR alert class
- seattle 12/9/03 5:45:41 PM PST Replication from Primary RVG hr\_rvg in dg hrdg on host seattle to Secondary london has stopped. VVR alert class
- seattle 12/9/03 5:45:40 PM PST Replication from Primary RVG hr\_rvg in dg hrdg on host seattle to Secondary london has stopped. VVR alert class

## Information Displayed in the Summary View of the RDS List View

RDS Count	Displays the number of RDSs on the selected host.
RDS Name	Displays the name of the RDS and the RDS icon.
Primary	Displays the name of the Primary host.
Primary RVG State	Displays the state of the Primary RVG.
Secondaries	Displays the number of Secondary hosts in the RDS.
Replication Status	Displays the status of replication between the Primary and the Secondaries in the RDS.

## Viewing RDS Information

The RDS view displays detailed information about the selected Replicated Data Set (RDS).

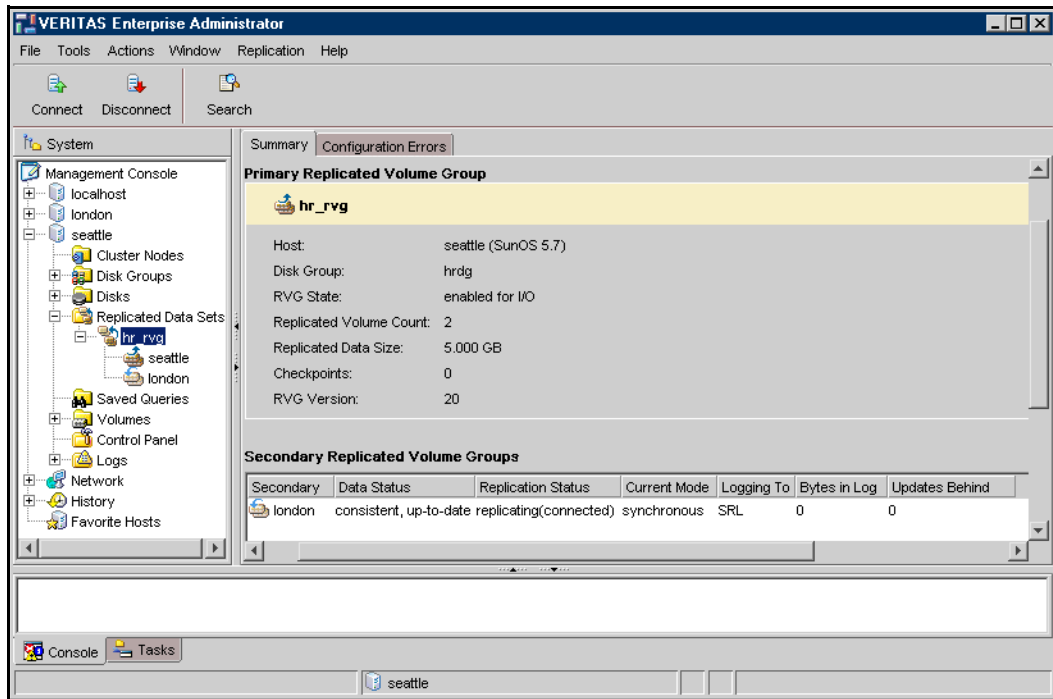
### Using the RDS View

Use the RDS view to get detailed information about a selected Replicated Data Set (RDS). To display the RDS view, expand the Replicated Data Sets node and select the RDS name in the tree view. For example, expand the Replicated Data Sets node in the tree view, and then click the RDS name **hr\_rvg** to display the RDS view for **hr\_rvg**.

By default, the Summary tab displays information about the selected RDS. Select the Configuration Errors tab in the RDS view to check whether configuration errors exist for the selected RDS.



## RDS View



## Information Displayed in the RDS View

The Summary tab in the RDS view displays the following high-level information about the Primary RVG in its top region:

Primary Replicated Volume Group:	
Host	The name of the host on which the Primary RVG is present and the version of the operating system on the host.
Disk Group	The name of the disk group.
RVG State	State of the RVG, whether enabled or disabled for I/O.
Replicated Volume Count	Number of data volumes in the RVG.
Replicated Data Size	The sum of the sizes of each data volume in the RVG.

Checkpoints	Number of checkpoints.
RVG Version	The version of the RVG.

- ◆ The table in the lower region of the RDS View (Summary tab) displays the following details for each Secondary in the selected RDS:

Secondary	Name of the Secondary host.
Data Status	Status of replication on the Secondary. For example, consistent, up-to-date. For more information about the values of the data status, refer to “ <a href="#">Displaying Consolidated Replication Status</a> ” on page 103.
Replication Status	Status of replication on the Secondary. For example, replicating (connected).
Current Mode	Displays the mode of replication, that is, synchronous or asynchronous.
Logging To	Indicates whether writes are being logged to the DCM or SRL.
Bytes in Log	The number of bytes in the SRL or DCM.
Updates Behind	Indicates the extent to which the Secondary is behind the Primary.

## Viewing Information About RVGs

When you double-click the RDS name or click the plus sign next to the RDS name in the tree view, the names of the Primary host and Secondary hosts in the RDS are displayed in the tree view. VVR VEA displays the names of the Primary and Secondary hosts and not the names of the Primary and Secondary RVG because the names of the Primary and Secondary RVG are the same on all the hosts in the RDS. Also, the name of the RDS will be the same as the name of the Primary or Secondary RVG.

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



## Using the RVG Views

### ▼ To display the RVG view, do one of the following:

- ✓ Expand the RDS name node, and then select the Primary or Secondary host name in the tree view.
- ✓ In the tree view, double-click the RDS name, and then select the Primary or Secondary host name.

## Viewing Information About the Primary RVG

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

- ✓ Expand the RDS name node, and select the Primary host name in the tree view.
- ✓ In the tree view, double-click the RDS name and select Primary host name.

The Primary RVG view is displayed. For a description of the information displayed in the Primary RVG View, see [“Information displayed in the Primary RVG View and the Secondary RVG View”](#) on page 272.

### Primary RVG View

The screenshot shows the VERITAS Enterprise Administrator interface. The left pane displays a tree view of the system hierarchy, with 'seattle' selected under 'Replicated Data Sets'. The main pane shows the 'Primary Replicated Volume Group View' for 'hr\_rvg'. The view includes a summary of the RVG and a table of replication attributes for remote hosts.

**Primary Replicated Volume Group**

**hr\_rvg**

Host: seattle (SunOS 5.7)  
 Disk Group: hrdg  
 RVG State: enabled for I/O  
 Replicated Volume Count: 2  
 Replicated Data Size: 5,000 GB  
 Checkpoints: 0  
 RVG Version: 20

**Replication attributes to remote hosts**

Secondary	Replication IP	Replication Mode	SRL Protection	Latency Protection	Latency Low Mai
london	london (192.168.225.159)	synchronous-override	autodcm	off	9950


## Viewing Information About the Secondary RVG

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

- ✓ Expand the RDS name node, and select the Secondary host name in the tree view.
- ✓ In the tree view, double-click the RDS name and select Secondary host name.

The Secondary RVG view is identical to the Primary RVG view, except that the Secondary RVG view does not contain the **Checkpoints** tab in the top region and the Secondary column in the lower region. The Secondary RVG view displays the same information as the Primary RVG view. See [“Information displayed in the Primary RVG View and the Secondary RVG View”](#) on page 272.

### Secondary RVG View



The screenshot shows the VERITAS Enterprise Administrator interface. The left pane displays a tree view with the following structure:

- System
  - Management Console
  - localhost
  - london
  - seattle
    - Cluster Nodes
    - Disk Groups
    - Disks
    - Replicated Data Set
      - hr\_rvg
      - seattle
      - london
    - Saved Queries
    - Volumes
    - Control Panel
    - Logs
  - Network
  - History
  - Favorite Hosts

The main pane displays the **Secondary Replicated Volume Group View** for the **hr\_rvg**. The view includes the following information:

**Secondary Replicated Volume Group**

- Host: london (SunOS 5.7)
- Disk Group: hrdg
- Replicated Volume Count: 2
- Replicated Data Size: 5,000 GB
- RVG Version: 20

**Replication attributes to remote hosts**

Replication IP	In Use	Replication Mode	SRL Protection	Latency Protection	Latency Low Mark	Later
seattle (192.168.225.158)	Yes	synchronous-override	autodcm	off	9950	10000



### Information displayed in the Primary RVG View and the Secondary RVG View

The layout for RVG views is as follows.

- ◆ The upper region contains tabs in which related information about the selected RVG is displayed. The **Summary** tab displays information that is identical to that of the RDS view. For details, see [“Information Displayed in the RDS View”](#) on page 268.
- ◆ The table in the lower region of RVG View displays the attributes of replication to the remote hosts.

The information about the selected RVG that is displayed in the various tabs is as follows:

Summary tab	<p>Displays the following information about the RVG:</p> <p>Host: Name of the host of the RVG. In the Primary RVG View, host name is the name of the Primary host. In the Secondary RVG View, the host name is name of the Secondary host.</p> <p>Disk Group: Name of the disk group.</p> <p>Replicated Volume Count: Number of data volumes in the RVG</p> <p>Replicated Data Size: The sum of the sizes of the data volumes in the RVG.</p> <p>Checkpoint: Number of checkpoints.</p> <p>Version: RVG version.</p> <p>Note that the type of the view, Primary RVG view or Secondary RVG view, and the RVG name is displayed at the top of the view.</p>
Data Volumes	<p>Displays information about the data volumes associated with the selected Primary or Secondary RVG. Information includes:</p> <p>Replicated Volume Count: Number of data volumes in the RVG.</p> <p>Replicated Data Size: The sum of the sizes of the data volumes in the RVG.</p> <p>Volume Name: Name of the volume</p> <p>Size: Size of the volume</p> <p>Layout: Layout of the volume</p> <p>Mirrors: Number of mirrors.</p> <p>DCM: Indicates whether or not the volume has a Data Change Map (DCM) associated with it.</p>



SRL	Displays the following information about the SRL: Name: Name of the SRL Size: Size of the SRL Layout: Layout of the SRL Mirrors: Number of mirrors
Checkpoints	In the Primary RVG view, the checkpoints are listed. If checkpoints do not exist, VVR VEA indicates that checkpoints are not found. In the Secondary RVG view, this tab does not display. To get the latest checkpoints, click <b>Get Checkpoints</b> .
Configuration Errors	Display problems with the configuration, if any.

The table in the lower region of the Summary tab in the RVG Views displays the replication attributes to the remote host.

Secondary	In the Primary RVG view, the name or names of the Secondary hosts are displayed. In the Secondary RVG view, this column is not displayed.
Replication IP	The host name and IP address used for replication.
Replication Mode	Displays the mode of replication, that is, synchronous or asynchronous.
SRL Protection	Displays autodcm, dcm, override, or off depending on the chosen option.
Latency Protection	Displays override, or off depending on the chosen option.
Latency-Low Mark	Number of updates in the SRL before the latency protection becomes inactive and updates succeed.
Latency-High Mark	Specifies the maximum number of waiting updates in the SRL before the latency protection becomes active and updates stall or fail.
Protocol	Transport Protocol being used for replication.
Bandwidth Limit	Displays the bandwidth limit that has been set for VVR. Displays <b>none</b> if a bandwidth limit has not been set, that is, Bandwidth Throttling is disabled.



## Administering Data Volumes

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

- ◆ [Associating a Data Volume to an RDS](#)
- ◆ [Resizing a Data Volume in a Replicated Data Set](#)
- ◆ [Dissociating a Data Volume from its RDS](#)

### Associating a Data Volume to an RDS

Use the Associate Volume task to associate a volume to an Replicated Data Set (RDS) even when replication is in progress. This command associates a volume to all the RVGs of the RDS.

Ensure that volumes of the same name and same length exist on the Primary and all the Secondaries of the RDS. You must create volumes of the required layout on the Primary and the Secondary.

---

**Caution** Failure to meet the following prerequisites may result in the Secondary data volumes having unusable data.

---

#### Prerequisites for associating a volume to an RDS

- ✓ Ensure that volumes of the same name and length as the Primary volumes exist on all the hosts in the RDS. If not, create them.
- ✓ Verify that the volume to be added is inactive.
- ✓ Synchronize the volume using the `vradm syncvol` command from the command-line interface (CLI).

---

**Note** By default, the Associate Volume task associates Data Change Maps (DCMs) to the volumes being added to the RDS, if DCMs have not been associated already. If any of the data volumes contains a DRL, this operation removes the DRL before associating the DCM to the volume.

---

#### ▼ To associate a volume with an RDS

1. Select the name of the RDS to which you want to associate the volume. For example, **hr\_rvg**.
2. Choose **Replication > Associate Volume**. To use the pop-up menu, right-click the name of the RDS in the tree-view and select **Associate Volume**.

3. In the Associate Volume dialog box, select the name of the data volume to be associated with the RDS. If the RVG contains VxVM ISP volumes, the list displays VxVM ISP volumes only.
4. Click **OK**. The result box displays a message, which indicates whether or not the operation succeeded.
5. In the result box, click **OK**.

## 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 while replication is in progress. The Resize Volume dialog box 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 resizes the file system also.
- ◆ 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, the original volume sizes are not restored. The volumes sizes will not 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.

### Prerequisite for resizing a data volume in an RDS

- ✓ Before increasing the size of a data volume, verify whether there is enough space in the disk group on the Primary and the Secondary.

### ▼ To resize a data volume in an RDS

1. From the tree view, double-click the RDS name and select the host name. For example, **seattle** or **london**.
2. From the Primary RVG view or Secondary RVG view, click the **Data Volume** tab to display the list of volumes.
3. Select the name of the data volume to be resized. For example, **hr\_dv01**.



4. Choose **Replication > Resize Volume**. To use the pop-up menu, right-click the name of the data volume and select **Resize Volume**.
5. In the **New Size** box, type the new size of the data volume. Note that this value must be an integer. You cannot enter decimal numbers.
6. Select the unit of measurement. The default is megabytes (MB).
7. To decrease the size of the data volume **hr\_dv01** when the size of the Primary and Secondary data volumes is the same, click **Force Resize**.
8. Click **OK**. The result box displays a message, which indicates whether or not the operation succeeded.
9. In the result box, click **OK**.

**Note:**

- ◆ Use the **Data Volume** tab in the RVG View to view the new size of the resized data volume.

## Dissociating a Data Volume from its RDS

Use the Dissociate Volume dialog box to remove a data volume from an RDS. The Dissociate Volume task only dissociates the data volume from the RVGs in an RDS; it does not delete the volume.

---

**Note** If you want the data volumes on the Primary and the Secondary hosts to be consistent after they are dissociated, stop the application, and make sure that the Secondaries are up-to-date and consistent before dissociating a data volume.

---

To remove a data volume from its RDS even when the application has not been stopped, select **Force Dissociate** from the **Dissociate Volume** dialog box.

### ▼ To dissociate a data volume from its RDS

1. Select the name of the RDS from which you want to dissociate the volume. For example, **hr\_rvg**.
2. Choose **Replication > Dissociate Volume**. To use the pop-up menu, right-click the name of the RDS in the tree view and then select **Dissociate Volume**. The Dissociate Volume dialog box is displayed.

3. From the drop-down box, select the name of the data volume to be dissociated from the RDS.
4. To dissociate the volume, if the application has not been stopped, click **Force Dissociate**.
5. Click **OK**. The result box displays a message, which indicates whether or not the operation succeeded.
6. In the result box, click **OK**.

## Increasing the Size of the SRL on the Primary

Use the Resize SRL task to increase the size of the Storage Replicator Log (SRL) on the Primary in a Replicated Data Set (RDS). The SRL Resize task enables you to increase the size of the Storage Replicator Log (SRL) on the Primary even while the application is active or while replication is in progress. The checkpoint information is preserved during this operation. The Resize SRL task does not increase the size of the SRL on the Secondary.

### Prerequisite for increasing the size of the SRL on the Primary in an RDS

- ✓ Before increasing the size of the SRL on the Primary, make sure that there is enough free space in the disk group in which the SRL resides.

### ▼ To increase the size of the SRL on the Primary in an RDS

1. From the tree view, double-click the RDS name and select the name of the Primary host. For example, **seattle**.
2. Choose **Replication > Resize SRL**. To use the pop-up menu, right-click the name of the Primary host and select **Resize SRL**.
3. In the **New Size** box, type the new size of the SRL. Note that this value must be an integer. You cannot enter decimal numbers.
4. Select the unit of measurement. The default is megabytes (MB).
5. Click **OK**. The result box displays a message, which indicates whether or not the operation succeeded.



6. In the result box, click **OK**.

**Note:**

- ◆ Use the **SRL** tab in the RVG View to view the new size of the resized SRL.
- ◆ For instructions on decreasing the size of the SRL on the Primary, see [“Decreasing the Size of the SRL on the Primary”](#) on page 404.
- ◆ For instructions on resizing the SRL on the Secondary, see [“Resizing the SRL on the Secondary”](#) on page 142.

## Administering Replication

VVR VEA enables you to administer replication using the following tasks:

- ◆ [Changing the Replication Settings](#)
- ◆ [Pausing Replication to a Secondary](#)
- ◆ [Resuming Replication to a Secondary](#)
- ◆ [Stopping Replication to a Secondary](#)
- ◆ [Changing the Replication Network for a Secondary](#)

## Changing the Replication Settings

The Change Replication Settings task enables you to set up or change the replication settings for the selected Secondary in an RDS. You can change the following replication settings:

- ◆ Mode of Replication
- ◆ Latency Protection
- ◆ SRL Overflow Protection
- ◆ Network Transport Protocol
- ◆ Packet Size
- ◆ Bandwidth Limit

Each replication setting affects performance and must be set up with care.

### ▼ To change the replication settings for a Secondary

1. Select the Secondary host name from the tree view, and then choose **Replication > Change Replication Settings**. To use the pop-up menu, right-click the name of the Secondary from the tree view, and then select **Change Replication Settings**.
2. Complete the Change Replication Settings dialog box. For instructions on how to complete the box, see [step 12](#) on page 255.
3. Click **OK**.

## Pausing Replication to a Secondary

Pausing replication to a Secondary prevents new and already-queued updates on the Primary from reaching the Secondary, and pauses communication between the Primary and Secondary.

Use the Pause Replication task to pause replication to a specific Secondary in an RDS. The Pause Replication task temporarily pauses replication between the Primary and Secondary. Pausing replication from the Primary enables you to perform administrative tasks, such as making changes to the network connecting the Primary and Secondary hosts. The Pause Replication task can be used when other applications on the host are required to use the entire network bandwidth.

### ▼ To pause replication to a Secondary

1. Select the Secondary host name from the tree view, and then choose **Replication > Pause Replication**. To use the pop-up menu, right-click the name of the Secondary, and then select **Pause Replication**.
2. In the Pause Replication box, click **OK**. A message indicates whether the pause replication task succeeded or failed.

**Notes:**

- ◆ When replication is paused, the RDS View displays the replication status for the Secondary RVG as **paused by user (primary paused)**.



## Resuming Replication to a Secondary

The Resume Replication task enables you to resume replication that was paused between the Primary and a Secondary.

### ▼ To resume replication to a Secondary

1. Select the Secondary host name from the tree view, and then choose **Replication > Resume Replication**. To use the pop-up menu, right-click the name of the Secondary, and then select **Resume Replication**.
2. Click **OK**. A message indicates whether the resume replication task succeeded or failed.

<b>Note:</b>
◆ The RDS View displays the replication status for the Secondary RVG as <b>replicating (connected)</b> after replication resumes.



## Stopping Replication to a Secondary

The Stop Replication task enables you to stop replication to a Secondary in an RDS.

---

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

---



**▼ To stop replication to a Secondary**

1. Select the Secondary host name from the tree view, and then choose **Replication > Stop Replication**. To use the pop-up menu, right-click the name of the Secondary.
2. In the Stop Replication box, click **OK**. A message indicates whether the stop replication task succeeded or failed.

**Notes:**

- ◆ The Stop Replication task fails if the Primary and Secondary are not up-to-date.
- ◆ To stop replication even when the Primary and Secondary are not up-to-date, select **Force Stop** in the Stop Replication dialog box.
- ◆ When replication is stopped, the RDS View displays the replication status for the Secondary RVG as **stopped (primary detached)**.



## Changing the Replication Network for a Secondary

You may need to change the host name or IP address of the Primary or the Secondary used for replication if you move the Primary or Secondary to a new location or if you want to use a different network for replication.

You can change the host name or IP address even after replication is established. The Change Replication Network task enables you to change the replication network between the Primary and a Secondary in an RDS. Note that you can change the replication network for either the Primary or the Secondary or both.

### ▼ To change the replication network for a Secondary

1. Select the Secondary host name from the tree view.
2. Choose **Replication > Change Replication Network**. To use the pop-up menu, right-click the name of the Secondary.
3. Complete the Change Replication Network dialog box as follows:

Primary	
Current	The current or existing name of the Primary is displayed.
New	Enter a new host name or IP address for the Primary, which can be used for replication. The new Primary host name or IP address must be resolvable and reachable from the Secondary host.
Secondary	
Current	The current or existing name of the Secondary is displayed.
New	Enter a new host name or IP address for the Secondary, which can be used for replication. This new Secondary host name or IP address must be resolvable and reachable from the Primary host.

---

**4. Click OK.****Notes:**

- ◆ If the Secondary name is changed, then the column **Replication IP** in the table in the Primary RVG View is updated with the new name.
- ◆ If the Primary name is changed, then the column **Replication IP** in the table in the Secondary RVG View is updated with the new name.



## Administering the Replicated Data Set

As part of administering an RDS, you may be required to remove a Secondary or a Primary from an RDS. This section describes how to perform these tasks:

- ◆ [Removing a Secondary from an RDS](#)
- ◆ [Removing a Primary RVG](#)

### Removing a Secondary from an RDS

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

---

**Caution** The Remove Secondary task is irreversible.

---

#### ▼ To remove a Secondary

1. Select the Secondary host name from the tree view, and then choose **Replication > Remove Secondary**. To use the pop-up menu, right-click the name of the Secondary, and then select **Remove Secondary**.
2. Click **OK**. A message indicates whether the Remove Secondary task succeeded or failed.

<b>Notes:</b>
<ul style="list-style-type: none"><li>◆ Replication must be stopped to the Secondary that is being removed. The Remove Secondary task fails if replication to the Secondary is not stopped.</li><li>◆ The Remove Secondary task does not delete data volumes; it only dissociates them from the Secondary.</li><li>◆ The Replicated Data Set View displays the updated RDS configuration.</li></ul>

## Removing a Primary RVG

The Remove Primary task removes a Primary RVG from an RDS and thus deletes the corresponding RDS.

### Prerequisite for deleting a Primary RVG:

- ✓ Make sure that all the Secondary RVGs have been removed from the RDS before performing the Remove Primary task. For instructions, see [“Removing a Secondary from an RDS”](#) on page 284.

---

**Note** This task can only be performed from the Primary host.

---

### ▼ To remove a Primary RVG from its RDS

1. Select the Primary host name from the tree view, and then choose **Replication > Remove Primary**. To use the pop-up menu, right-click the name of the Primary in the tree view, and then select **Remove Primary**.
2. To remove the Primary RVG while the application using the Primary data volumes is active, select **Force Remove**.
3. Click **OK**. A message indicates whether the Remove Primary task succeeded or failed.

<b>Note:</b>
◆ The Remove Primary task does not delete data volumes or the SRL from the VERITAS Volume Manager configuration.



## Incrementally Synchronizing the Secondary After SRL Overflow

If the SRL overflows and log overflow protection is set to **autodcm** or **dcm**, the Secondary can be incrementally synchronized using the Replay DCM task. After the Primary and all the Secondaries reconnect, the DCM resynchronization process must be started to transfer data from the Primary to the Secondary. The Replay DCM task enables you to replay an active DCM and incrementally resynchronize the Secondary after SRL overflow.

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 resynchronization, the applications cannot be restarted on the Secondary.



If the Secondary volumes are mirrored, you can break off mirrors to retain consistent (though out-of-date) copies of data until the resynchronization is complete. Use the snapshot feature to do this if snapshots exist.

▼ **To incrementally resynchronize the Secondary after SRL overflow**

1. From the tree view, select the name of the RDS for which you want to resynchronize the Secondary RVGs. For example, **hr\_rvg**.
2. Choose **Replication > Replay DCM**. To use the pop-up menu, right-click the name of the RDS.
3. In the Replay DCM dialog box, click **OK**. A message indicates whether the resynchronization process succeeded or failed.

**Notes:**

- ◆ Data is transmitted to the Secondaries only after all the Secondaries are connected to the Primary.
- ◆ All the Secondaries taking part in the resynchronization must remain connected for the resynchronization to continue. The resynchronization pauses if any of the Secondaries disconnects or is manually paused.
- ◆ While resynchronization is in progress, the RDS View displays the value **DCM** in the Logging To column.

For details about each method, refer to “[Transferring the Primary Role](#)” on page 215.

## 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 maintenance or to make the application active on another node.

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

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

For details about migrating the Primary role, refer to “[Transferring the Primary Role](#)” on page 215.

### ▼ To migrate a healthy Primary

1. Stop all the applications that are involved in replication. For example, if the application is a file system, unmount it.
2. Verify that the Secondaries are up to date by checking the RDS view. In the RDS View, the Data Status of the Secondary must be **consistent, up to date** and the Replication Status must be **connected (replicating)**.
3. From the tree view, select the name of Secondary host whose role you want to change. For example, **london**.
4. Choose **Replication > Migrate**. To use the pop-up menu, right-click the name of the RDS.



5. In the Migrate Primary Role dialog box, click **OK**. A message that indicates whether the migration succeeded or failed is displayed.

**Notes:**

- ◆ The Migrate task fails if the active Secondaries are not **up to date** and if the data volumes are active.
- ◆ By default, the Migrate task starts replication from the new Primary to the original Primary.
- ◆ If an RDS contains multiple Secondaries, VVR VEA does not automatically reconfigure the additional Secondaries.
- ◆ If required, you can reconfigure the additional Secondaries as Secondary hosts of the new Primary using the VVR command-line interface.
- ◆ The Data Status of the Secondary (attached RLINKs) must be **consistent, up to date** and the Replication Status must be **connected (replicating)**.
- ◆ If the Migrate task fails on any of the hosts, the original configuration is restored.
- ◆ On migration, the Secondary becomes the new Primary.
- ◆ Because the application was stopped before the migration, an application recovery is not required. To start the application before enabling replication after migration, first, pause replication using the Pause task (**Replication > Pause**), start the application, and then resume replication using the Resume task (**Replication > Resume**).



## 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 task 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 task performs the following functions on the RDS to which the original Primary belongs:

- ◆ Converts the Secondary RVG to a Primary RVG.
- ◆ Removes the original Primary from the RDS.
- ◆ Enables the fast-failback feature on the new Primary depending on the failback option that you select from the Takeover task. For information on failing back, see [“Failing Back to the Original Primary”](#) on page 292.

For details on taking over from the original Primary, refer to [“Transferring the Primary Role”](#) on page 215.

## Failback Options

The Takeover task 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 and is the recommended option. This option enables fast failback. Use this option if you want to incrementally synchronize the original Primary with the new Primary. However, conversion of the original Primary to a Secondary and synchronization of the original Primary data volumes with the data volumes on the new Primary is not done automatically.

To synchronize the data volumes on the original Primary and convert the original Primary to a Secondary, use the Replay Failback Log task (**Replication > Replay Failback Log**) on the RDS.

- ◆ **Auto fast failback**—Use this option if you want to convert the original Primary to a Secondary after the original Primary becomes available and also to automatically synchronize the data volumes on original Primary, which is now the new Secondary, using fast failback.



- ◆ **No fast failback** —If you choose this option, you must synchronize the original Primary after it becomes available. Use 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.

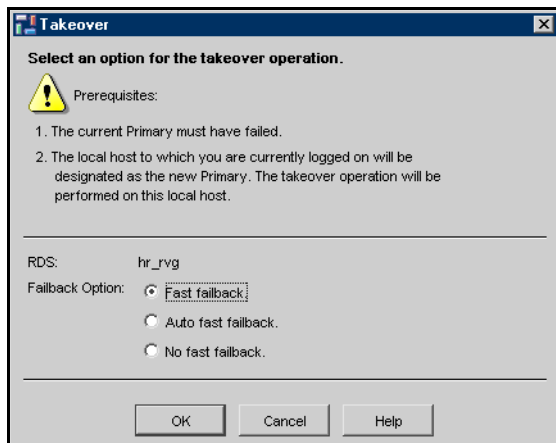
---

**Tip** We recommend that you use the **Fast failback** option to synchronize the original Primary with the new Primary. For instructions, see [“Failing Back Using Fast-Failback Synchronization”](#) on page 293.

---

▼ **To take over from the original Primary to the Secondary**

1. Make sure you have logged on to the Secondary host. To log on to a host from VVR VEA, choose **File > Connect** and specify the name of the Secondary host to which you want to connect.
2. From the tree view, select the name of Secondary host whose role you want to change to Primary. For example, **london**.
3. Choose **Replication > Takeover**. To use the pop-up menu, right-click the name of the RDS and select **Takeover**.



4. Complete the Takeover dialog box as follows:

<p>Options: Fast Failback Auto fast failback No fast failback</p>	<p>By default, <b>Fast failback</b> is selected. If this option is not appropriate, select the required option.</p> <ul style="list-style-type: none"> <li>◆ To enable fast failback logging only, select <b>Fast failback</b>. You can synchronize the data volumes on the original Primary and convert the original Primary to a Secondary using the Replay Failback Log task (<b>Replication &gt; Replay Failback Log</b>).</li> <li>◆ 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 <b>Auto fast failback</b>.</li> <li>◆ To change the role from Secondary to Primary without enabling fast failback, select <b>No fast failback</b>. 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.</li> </ul>
---	---

5. Click **OK**. A message that indicates whether the takeover succeeded or failed is displayed.
6. 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.

**Notes:**

- ◆ If you want to use the fast-failback feature, make sure that the data volumes on the Secondary host have associated DCMs.
- ◆ If fast failback is enabled, the RDS View displays the value **DCM** in the Logging To column for the original Primary.



## Failing Back to the Original Primary

After an unexpected failure, the original 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*.

VVR VEA provides the following methods to fail back to the original Primary:

- ◆ [Failing Back Using Fast-Failback Synchronization](#)
- ◆ [Failing Back Using Automatic Synchronization](#)

---

**Note** It is recommended that you use the fast-failback synchronization method.

---

The fast-failback synchronization method can be used when the Fast failback option is selected during Takeover.

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, you must synchronize the Secondary completely. In this case, you must failback using automatic synchronization. For details about failing back to the original Primary, refer to “[Failing Back to the Original Primary](#)” on page 292.

## Failing Back Using Fast-Failback Synchronization

To use the fast-failback synchronization method to fail back to the original Primary, the fast-failback feature must be enabled when taking over the Primary role.

For more information about the fast-failback feature, see the Chapter 7, “Transferring the Primary Role” in the *VERITAS Volume Replicator Administrator’s Guide*.

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

1. If **Auto fast failback** is selected in the Takeover dialog box during the takeover, go to [step 3](#); otherwise go to the next step.
2. If **Fast failback** was selected in the Takeover dialog box during the takeover, 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 task (**Replication > Replay Failback Log**).
3. After the replay completes, stop the application on the new Primary at a scheduled time.
4. Migrate the Primary Role to the original or failed Primary. For instructions, see “[Migrating the Primary Role](#)” on page 287.



## 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.

This example explains how to failback when fast failback is enabled on `london` by selecting **Auto fast failback** in the Takeover wizard or when **Fast failback** was selected.

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 using fast failback

1. If **Auto fast failback** was selected in the Takeover task when taking over from the original Primary `seattle`, go to [step 5](#); If **Fast failback** was selected, go to the next step.
2. From the tree view, select the name of the RDS for which you want to replay the failback log. For example, `hr_rvg`.
3. Choose **Replication > Replay Failback Log**. To use the pop-up menu, right-click the name of the RDS.
4. Click **OK**. When the synchronization completes, go to the next step. You can check the status of the synchronization using the RDS view.

The Replay Failback Log task 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 task when taking over from the original Primary.

5. At a convenient time, after verifying that the synchronization has completed, stop the application on the new Primary `london`.
6. Migrate the Primary role from the new Primary host `london` to the original Primary host `seattle` using the Migrate task (**Replication > Migrate**). For instructions, see [“Migrating the Primary Role”](#) on page 287.

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

## Failing Back Using Automatic Synchronization

### ▼ To fail back to the original Primary using automatic synchronization

1. Convert the original Primary to a Secondary using the **Make Secondary** task (**Replication > Make Secondary**).

For detailed instructions, see [“Converting a Primary to a Secondary”](#) on page 296.

2. Synchronize the data volumes on the original Primary with the data volumes on the new Primary using the Synchronize RVG task (**Replication > Synchronize RVG**).

For detailed instructions, see [“Synchronize the Secondary RVG with the Primary RVG”](#) on page 297.

3. Stop the application on the new Primary.

4. Migrate the Primary role from the new Primary host to the original Primary host. Replication from the original Primary to the original Secondary is started by default. See [“Migrating the Primary Role”](#) on page 287.



## Converting a Primary to a Secondary

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

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

---

**Note** Use the Make Secondary task to fail back to the original Primary only if **Fast failback** was not selected during the takeover from the original Primary.

---

### ▼ To convert an original Primary to a Secondary

1. From the tree view, select the name of the original Primary host that has restarted and whose role you want to convert to Secondary. For example, **seattle**.
2. Choose **Replication > Make Secondary**. To use the pop-up menu, right-click the name of the RDS.
3. Complete the Make Secondary dialog box as follows:

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

4. Click **OK**. A message that indicates whether the Make Secondary task succeeded or failed is displayed.

**Notes:**

- ◆ Before using the Make Secondary task, make sure that you stop all the applications that use the data volumes on the original Primary. Also, ensure that none of the data volumes are open.



## Synchronize the Secondary RVG with the Primary RVG

Use the Synchronize RVG task to synchronize the data volumes on the Secondary with the data volumes on the Primary. The Synchronize RVG task can be used in the failback procedure to synchronize the data volumes on the original Primary with the data volumes on the new Primary using automatic synchronization.

### ▼ To synchronize the Secondary RVG with the Primary RVG

1. Select the Secondary host name from the tree view.
2. Choose **Replication > Synchronize RVG**.
3. Click **OK**. The status of the synchronization is displayed in a message box.

#### Note:

- ◆ While resynchronization is in progress, the RDS View displays the values **resync in progress** and **DCM (autosync)** in the Replication Status and Logging To columns respectively.

## Example: Failing Back Using Automatic Synchronization

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

### ▼ To fail back using automatic synchronization

1. From the tree view, select the name of original Primary host, **seattle**, which has restarted and whose role you want to convert to Secondary.
2. Choose **Replication > Make Secondary**. To use the pop-up menu, right-click the name of the RDS.
3. In the Make Secondary box, select or enter **london** in the **New Primary** box. For detailed instructions, see [“Converting a Primary to a Secondary”](#) on page 296.

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

4. Select the Secondary host name **seattle** from the tree view.



5. Choose **Replication > Synchronize RVG**. For detailed instructions, see [“Synchronize the Secondary RVG with the Primary RVG”](#) on page 297.

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 automatic synchronization.

6. Stop the applications on the new Primary **london**.
7. Migrate the Primary role from the new Primary host **london** to the original Primary host **seattle** using the Migrate wizard (**Replication > Migrate**). For instructions, see [“Migrating the Primary Role”](#) on page 287. Replication from the original Primary **seattle** to the original Secondary **london** is started by default after the migration is completed.

## Event Notification Using VVR VEA

VERITAS Volume Replicator (VVR) provides event notification by SMTP email and through SNMP traps that can be displayed in HP OpenView, CA Unicenter, and IBM Tivoli. Note that the notification service is available only when VVR VEA is installed. You can configure the notification service to send messages to specific individuals or groups and also set up rules for filtering the event messages. You can also set up rules that will run a script when specified conditions apply.

This section presents an overview of the Rule Manager utility, which provides the event notification service for VVR. The full details for this utility are given in the online help module titled “Optional Rules for Handling Events.”

---

**Note** VVR sends SNMP traps to port 162 on the SNMP manager.

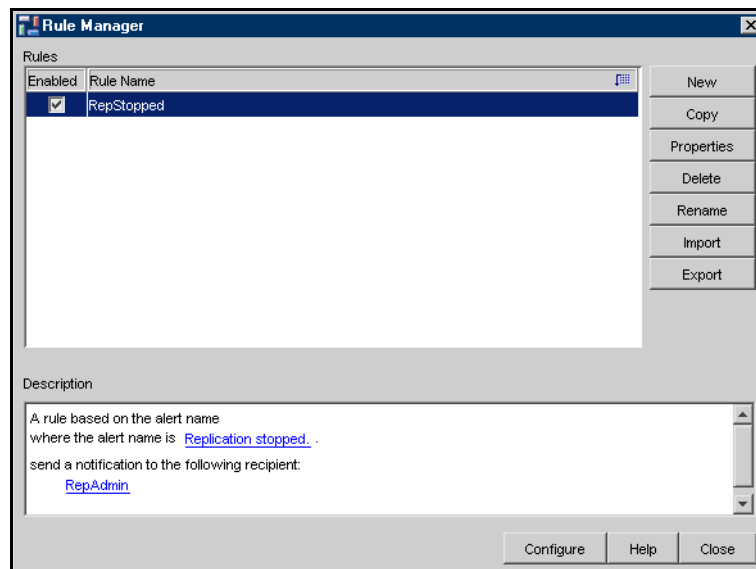
---

### Accessing the Rule Manager Utility

The commands for configuring the service are all accessed through the **Rule Manager** icon in the VEA Control Panel. If you click the Control Panel node in the tree view, the Control Panel icons display in the right pane of the VEA window.

Double-click the **Rule Manager** icon to bring up the Rule Manager window. All functions relating to event notification rules are initiated from this window.

Rule Manager Window



From the Rule Manager, you can create, enable, disable, modify, and delete rules relating to event notification. The **Configure** button in the lower portion of the window is used to configure settings, such as the settings for the SMTP mail server and the recipient list of people or groups to be notified when the conditions of a specified rule are met.

## Preliminary Setup

Before creating rules, you must provide information required for automated actions.

### ▼ To provide settings for automated actions

1. Configure the SMTP mail notification settings you want to use. To accomplish this, do the following:
  - a. From the VERITAS Enterprise Administrator Control Panel, select **Rule Manager**.  
The Rule Manager window is displayed.
  - b. Click **Configure**; then click **Configure SMTP**.
  - c. Provide information about your SMTP mail server, port, and default sender.  
While working in the dialog box, select **Help** at any time for details about the information you must provide.

This is a one-time setting. After providing this information, you don't need to repeat this step unless the SMTP mail server changes for your system.

2. Set up a list of recipients to be available for notification actions in rules. To accomplish this, do the following:
  - a. From the VERITAS Enterprise Administrator Control Panel, select **Rule Manager**.  
The Rule Manager window is displayed.
  - b. Click **Configure**; then click **Define Recipients**.

- c.** Select **New** and define a recipient or group of recipients.

Repeat this step for each recipient or group you want to add to the list. You have the following options for defining recipients:

- ◆ Temporarily disable notification for a recipient.
- ◆ Limit when notifications can be sent to a recipient.
- ◆ Limit how many notifications can be sent to a recipient during a given period of time.

While defining recipients, select **Help** at any time for details about the fields and elements in the active dialog box.

- 3.** If you are planning to send SNMP traps, the MIB files should be loaded and compiled into the management application to which traps will be sent. To accomplish this, do the following:

- a.** Load the product-specific MIB into your SNMP management software database according to instructions provided by your management platform.

For details, see “[Accessing the VERITAS Volume Replicator SNMP MIB Files](#)” on page 302.

- b.** If the platform you are using requires a configuration file, such as `trapd.conf`, you should load that file as well. Consult the user’s guide for the management platform to determine how to map your traps to the alarm system of the platform.

Any configuration files associated with your product are stored in the same location as the MIB files. VVR sends SNMP v1 traps.



## Main Steps for Creating Rules

This section explains how to create rules.

### ▼ To create rules

1. From the VEA Control Panel, select **Rule Manager**.  
The Rule Manager window is displayed.
2. Click **New** to start the Rule wizard.
3. Follow the wizard prompts to do the following:
  - a. Select the type of rule you want to create. You have two options:
    - ◆ Create a rule for certain alerts that you identify by name.
    - ◆ Create a rule for all alerts that have a particular severity and/or classification.
  - b. Choose the events to monitor.
  - c. Configure one or more actions to be taken when the events are detected.
  - d. Provide a name and description for the rule.

---

**Note** For details on creating rules, see the “Optional Rules for Handling Events” online help module. To access this module, select **Contents** from the Help menu.

---

## Accessing the VERITAS Volume Replicator SNMP MIB Files

If you want to receive VVR SNMP traps on a management platform (such as HP OpenView), you will need to compile and load the VERITAS SNMP MIB files into your management platform, following the procedures outlined in the documentation for your specific platform.

When you install VVR, the VERITAS MIB files are installed in the following directory:

```
/opt/VRTSob/sig/snmpmibfiles
```

The VVR MIB file is **VRTS-VVR-UX.mib**.

This chapter describes the different kinds of configuration errors that could occur when using VVR and explains how to correct these errors.

## Recovery from RLINK Connect Problems

This section describes the errors that may be encountered when connecting RLINKs. To be able to troubleshoot RLINK connect problems, it is important to understand the RLINK connection process.

Connecting the Primary and Secondary RLINKs is a two-step operation. The first step, which attaches the RLINK, is performed by issuing the `vxlink att` command or the `vradmin startrep` command. The second step, which connects the RLINKs, is performed by the kernels on the Primary and Secondary hosts.

The `attach` command performs a number of checks to ensure that the operation is likely to succeed, and if it does, the command changes the state of the RLINKs from `DETACHED/STALE` to `ENABLED/ACTIVE`. The command then returns success.

If the command is successful, the kernel on the Primary is notified that the RLINK is enabled and it begins to send messages to the Secondary requesting it to connect. Under normal circumstances, the Secondary receives this message and connects. The state of the RLINKs then changes from `ENABLED/ACTIVE` to `CONNECT/ACTIVE`.

If the RLINK does not change to the `CONNECT/ACTIVE` state within a short time, there is a problem preventing the connection. This section describes a number of possible causes. An error message indicating the problem may be displayed on the console.



- ◆ If the following error displays on the console:

```
VxVM VVR vxrlink INFO V-5-1-5298 Unable to establish connection
with remote host <remote_host>, retrying
```

Make sure that the `vradmin` daemon is running on the Primary and the Secondary hosts; otherwise, start the `vradmin` daemon by issuing the following command:

```
# /sbin/init.d/vras-vradmin.sh start
```

- ◆ If the following error displays on the console:

```
VxVM VVR vxrlink ERROR V-5-1-3504 Primary datavol hr_dv01 is not
mapped on secondary, yet
```

Make sure that a corresponding Secondary volume mapped to the Primary data volume `hr_dv01` exists on the Secondary. If the volume has a different name, either rename the volume, or map the name of the volume on the Primary with the name of the volume on the Secondary. To map the volume on the Secondary with the volume on the Primary, issue the following command on the Secondary:

```
# vxedit -g hrdg set primary_datavol=hr_dv01 hr_dv01_sec
```

where `hr_dv01_sec` is the name of the volume on the Secondary.

- ◆ If there is no self-explanatory error message, issue the following command on both the Primary and Secondary hosts:

```
# vxprint -g diskgroup -l rlink_name
```

In the output, check the following:

The `remote_host` of each host is the same as `local_host` of the other host.

The `remote_dg` of each host is the same as the disk group of the RVG on the other host.

The `remote_dg_dgid` of each host is the same as the `dgid` (disk group ID) of the RVG on the other host as displayed in the output of the `vxprint -l diskgroup` command.

The `remote_rlink` of each host is the same as the name of the corresponding RLINK on the other host.

The `remote_rlink_rid` of each host is the same as the `rid` of the corresponding RLINK on the other host.

Make sure that the network is working as expected. Network problems might affect VVR, such as prevention of RLINKs from connecting or low performance. Possible problems could be high latency, low bandwidth, high collision counts, and excessive dropped packets.



- ◆ For an RLINK in private disk group, issue the following command on each host.

```
# ping <remote_host>
```

There should be no packet loss or very little packet loss. To ensure that the network can transmit large packets issue the following command on each host for an RLINK in a private disk group.

```
# ping remote_host 8192
```

The packet loss should be about the same as for the earlier ping command.

- ◆ Issue the vxiod command on each host to ensure that there are active I/O daemons. If the output is 0 volume I/O daemons running, activate I/O daemons by issuing the following command:

```
# vxiod set 10
```

- ◆ VVR uses well-known ports to establish communications with other hosts.

Issue the following command to display the port number:

```
# vxprint -g diskgroup -l rlink_name
```

Issue the following command to ensure that the heartbeat port number in the output matches the port displayed by vxprint command:

```
# vrport
```

Confirm that the heartbeat port has been opened by issuing the following command:

```
# netstat -an | grep port-number
```

where port-number is the port number being used by the heartbeat server as displayed by the vrport command.

The output looks similar to this:

```
udp      0      0 *.port-number
```



## Recovery from Configuration Errors

Configuration errors occur when the configuration of the Primary and Secondary RVGs is not identical. Each data volume in the Primary RVG must have a corresponding data volume in the Secondary RVG of exactly the same size; otherwise, replication will not proceed.

Errors in configuration are detected in two ways:

- ◆ When an RLINK is attached for the first time, the configuration of the Secondary is checked for configuration errors. If any errors are found, the `attach` command fails and prints error messages indicating the problem.
- ◆ Changes that affect the configuration on the Primary or Secondary may cause the Secondary to enter the PAUSE state with the `secondary_config_err` flag set. The problem is fixed by correcting the configuration error, and then resuming the RLINK.

### Errors During an RLINK Attach

When an RLINK is attached, VVR checks whether for each data volume associated to the Primary RVG, the Secondary RVG has an associated data volume of the same size that is mapped to its counterpart on the Primary. The following example illustrates an attempted attach with every possible problem and how to fix it. Before the `attach`, the Primary has this configuration:

TY	Name	Assoc	KSTATE	LENGTH	STATE
rv	hr_rvg	-	DISABLED	-	EMPTY
r1	r1k_london_hr_rvg	hr_rvg	DETACHED	-	STALE
v	hr_dv01	hr_rvg	ENABLED	12800	ACTIVE
pl	hr_dv01-01	hr_dv01	ENABLED	12800	ACTIVE
sd	disk01-05	hr_dv01-01	ENABLED	12800	-
v	hr_dv02	hr_rvg	ENABLED	12800	ACTIVE
pl	hr_dv02-01	hr_dv02	ENABLED	12880	ACTIVE
sd	disk01-06	hr_dv02-01	ENABLED	12880	
v	hr_dv03	hr_rvg	ENABLED	12880	ACTIVE
pl	hr_dv03-01	hr_dv03	ENABLED	12880	ACTIVE
sd	disk01-07	hr_dv03-01	ENABLED	12880	-

v	hr_srl	hr_rvg	ENABLED	12880	ACTIVE
pl	hr_srl-01	hr_srl	ENABLED	12880	ACTIVE
sd	disk01-08	hr_srl-01	ENABLED	12880	0 -

The Secondary has the following configuration:

TY	Name	Assoc	KSTATE	LENGTH	STATE
rv	hr_rvg	-	ENABLED	-	- ACTIVE
rl	rlk_seattle_hr_rvg	hr_rvg	ENABLED	-	- ACTIVE
v	hr_dv01	hr_rvg	ENABLED	12700	- ACTIVE
pl	hr_dv01-01	hr_dv01	ENABLED	13005	- ACTIVE
sd	disk01-17	hr_dv01-01	ENABLED	13005	0 -
v	hr_dv2	hr_rvg	ENABLED	12880	- ACTIVE
pl	hr_dv02-01	vol2	ENABLED	13005	- ACTIVE
sd	disk01-18	hr_dv02-01	ENABLED	13005	0 -
v	hr_srl	hr_rvg	ENABLED	12880	- ACTIVE
pl	hr_srl-01	hr_srl	ENABLED	13005	- ACTIVE
sd	disk01-19	hr_srl-01	ENABLED	13005	0 -



Note that on the Secondary, the size of volume `hr_dv01` is small, `hr_dv2` is misnamed (must be `hr_dv02`), and `hr_dv03` is missing. An attempt to attach the Primary RLINK to this Secondary using the `attach` command fails.

```
# vxrlink -g hrdg -f att rlk_london_hr_rvg
```

The following messages display:

```
VxVM VVR vxrlink INFO V-5-1-3614 Secondary data volumes detected
with rvg hr_rvgas parent:
VxVM VVR vxrlink ERROR V-5-1-0 Size of secondary datavol hr_dv01
(len=12700) does not match size of primary (len=12800)
VxVM VVR vxrlink ERROR V-5-1-3504 primary datavol hr_dv02 is not
mapped on secondary, yet
VxVM VVR vxrlink ERROR V-5-1-3504 primary datavol hr_dv03 is not
mapped on secondary, yet
```

To fix the problem, issue the following commands on the Secondary:

1. Resize the data volume `hr_dv01`:

```
# vradmin -g hrdg resizevol hr_rvg hr_dv01 12800
```

2. Rename the data volume `hr_dv2` to `hr_dv02`:

```
# vxedit -g hrdg rename hr_dv2 hr_dv02
```

3. Associate a new volume, `hr_dv03`, of the same size as the Primary data volume `hr_dv03`.

```
# vxassist -g hrdg make hr_dv03 12800
# vxvol -g hrdg assoc hr_rvg hr_dv03
```

Alternatively, the problem can be fixed by altering the Primary to match the Secondary, or any combination of the two. When the Primary and the Secondary match, retry the attach.

On the Primary:

```
# vxrlink -g hrdg -f att rlk_london_hr_rvg
```

```
VxVM VVR vxrlink INFO V-5-1-3614 Secondary data volumes detected
with rvg hr_rvg as parent:
VxVM VVR vxrlink INFO V-5-1-0 vol1: len=12800
primary_datavol=hr_dv01
VxVM VVR vxrlink INFO V-5-1-0 vol1: len=12800
primary_datavol=hr_dv02
VxVM VVR vxrlink INFO V-5-1-0 vol1: len=12800
primary_datavol=hr_dv03
```

## Errors During Modification of an RVG

After the initial setup and attach of a Secondary RLINK, incorrect modifications such as adding, resizing, and renaming volumes can cause the affected RLINK to be PAUSED with the `secondary_config_err` flag set. This prevents replication to the Secondary until the problem is corrected.

Run the `vxrlink verify rlink` command at either node to check whether this has occurred. When the configuration error has been corrected, the affected RLINK can be resumed.

### Missing Data Volume Error

If a data volume is added to the Primary RVG and the Secondary has no corresponding data volume, the RLINK state changes to PAUSED with the `secondary_config_err` flag set. Executing the `vxrlink verify` command produces the following:

On the Primary:

```
# vxrlink -g hrdg verify rlk_london_hr_rvg
```

RLINK	REMOTE HOST	LOCAL_HOST	STATUS	STATE
rlk_london_hr_rvg	london	seattle	ERROR	PAUSE

ERROR: hr\_dv04 does not exist on Secondary (london)

On the Secondary:

```
# vxrlink -g hrdg verify rlk_seattle_hr_rvg
```

RLINK	REMOTE HOST	LOCAL_HOST	STATUS	STATE
rlk_seattle_hr_rvg	seattle	london	ERROR	PAUSE

ERROR: hr\_dv04 does not exist on Secondary (local host)

To correct the problem, either create and associate `hr_dv04` on the Secondary or alternately, dissociate `vol04` from the Primary, and then resume the Secondary RLINK. To resume the Secondary RLINK, use the `vradm resumerrep rvg_name` command.

If `hr_dv04` on the Primary contains valid data, copy its contents to `hr_dv04` on the Secondary before associating the volume to the Secondary RVG.



## Data Volume Mismatch Error

If a Primary data volume is increased in size, but the Secondary data volume is not, a configuration error results.

On the Primary:

```
# vxassist growby hr_dv04 100
# vxrlink -g hrdg verify rlk_london_hr_rvg
```

RLINK	REMOTE_HOST	LOCAL_HOST	STATUS	STATE
rlk_london_hr_rvg	london	seattle	ERROR	PAUSE

ERROR: hr\_dv04 too small (12800 blocks). primary is 12900

On the Secondary:

```
# vxrlink -g hrdg verify rlk_seattle_hr_rvg
```

RLINK	REMOTE_HOST	LOCAL_HOST	STATUS	STATE
rlk_seattle_hr_rvg	seattle	london	ERROR	PAUSE

ERROR: hr\_dv04 too small (12800 blocks). primary is 12900

To correct the problem, increase the size of the Secondary data volume, or shrink the Primary data volume:

```
# vradmin -g hrdg resizevol hr_rvg hr_dv04 12900
```

After resizing a data volume, resume the Secondary RLINK by issuing the following command on any host in the RDS:

```
# vradmin -g hrdg resumerep hr_rvg
```



## Data Volume Name Mismatch Error

If a volume is renamed on the Primary but not on the Secondary, a configuration error results and the RLINK will be disconnected. Use the `vxprint -lP` command to view the RLINK flags. If the `secondary_config_err` flag is set, use one of the following commands to determine if there is a data volume name mismatch error.

On the Primary:

```
# vxrlink -g hrdg verify rlk_london_hr_rvg

RLINK          REMOTE_HOST    LOCAL_HOST    STATUS    STATE
rlk_london_hr_rvg  london        seattle      ERROR    PAUSE
ERROR: hr_dv04 on secondary has wrong primary_datavol name (hr_dv04,
should be hr_dv05)
```

On the Secondary:

```
# vxrlink -g hrdg verify rlk_seattle_hr_rvg

RLINK          REMOTE_HOST    LOCAL_HOST    STATUS    STATE
rlk_seattle_hr_rvg  seattle        london       ERROR    PAUSE
ERROR: hr_dv04 on secondary has wrong primary_datavol name (hr_dv04,
should be hr_dv05)
```

To fix this error, do one of the following:

- ◆ Rename either the Primary or Secondary data volume, and resume the RLINK using the `vradmin resumerep rvg_name` command.

OR

- ◆ Set the `primary_datavol` field on the Secondary data volume to refer to the new name of the Primary data volume as follows, and resume the RLINK using the `vradmin resumerep rvg_name` command.

On the Secondary:

```
# vxedit -g hrdg set primary_datavol=hr_dv05 hr_dv04
```

where `hr_dv05` is the new name on the Primary



## Recovery

### Primary-Host Crash

When a Primary host recovers from a failure, VVR automatically recovers the RVG configuration. When the Primary recovers, VVR recovers the Primary SRL and all volumes in the RVG. Information about the recent activity on the SRL and the data volume is maintained in the SRL header. VVR uses this information to speed up recovery, which is automatic on reboot.

### Recovering from Primary Data Volume Error

If there is an error during access of a Primary data volume, the data volume is detached, the RVG is disabled, and the RVG state changes to FAIL, which means that the RVG is unusable. RLINKs are not affected. If the SRL was not empty at the time of the volume error, those updates will continue to flow from the SRL to the Secondary RLINKs.

Recovery from this failure consists of two parts:

- ◆ Restoring the Primary data volume from backup
- ◆ Resynchronizing any Secondary RLINKs

If the RVG contains a database, recovery of the failed data volume must be coordinated with the recovery requirements of the database. The details of the database recovery sequence determine what must be done to synchronize Secondary RLINKs.

Two examples are given below.

#### Example 1

In this example, all the RLINKs are detached before recovery of the failure begins on the Primary. When recovery of the failure is complete, including any database recovery procedures, all the RLINKs must be synchronized using a Primary checkpoint.

On the Primary (*seattle*):

1. Detach all RLINKs

```
# vxrlink -g hrdg det rlk_london_hr_rvg
```



2. Fix or repair the data volume.

If the data volume can be repaired by repairing its underlying subdisks, you need not dissociate the data volume from the RVG. If the problem is fixed by dissociating the failed volume and associating a new one in its place, the dissociation and association must be done while the RVG is stopped.

3. Make sure the data volume is started before restarting the RVG.

```
# vxvol -g hrdg start hr_dv01
# vxrvrg -g hrdg start hr_rvg
```

4. Restore the database.

5. Synchronize all the RLINKs using block-level backup and checkpointing.

### Example 2

This example does the minimum to clear the RVG FAIL state leaving all RLINKs attached. In this example, restoring the failed volume data from backup, and the database recovery is done with live RLINKs. Because all the changes on the Primary are replicated, all the Secondaries must be consistent with the Primary after the changes have been replicated. This method may not always be practical because it might require replication of large amounts of data. The repaired data volume must also be carefully tested on every target database to be supported.

On the Primary (*seattle*):

1. Stop the RVG.

```
# vxrvrg -g hrdg stop hr_rvg
```

2. Dissociate the failed data volume from the RVG.

3. Fix or repair the data volume or use a new volume.

If the data volume can be repaired by repairing its underlying subdisks, you need not dissociate the data volume from the RVG. If the problem is fixed by dissociating the failed volume and associating a new one in its place, the dissociation and association must be done while the RVG is stopped.

4. Associate the volume with the RVG.

5. Make sure the data volume is started before restarting the RVG. If the data volume is not started, start the data volume:

```
# vxvol -g hrdg start hr_dv01
```



6. Start the RVG:

```
# vxrvrg -g hrdg start hr_rvg
```

7. Restore the database.

As an alternative to the procedures described in Example 1 and Example 2, the Primary role can be transferred to a Secondary host. For details, see Chapter 7, “[Transferring the Primary Role](#)” on page 215.

## Primary SRL Volume Error Cleanup and Restart

If there is an error accessing the Primary SRL, the SRL is dissociated and the RLINKs are detached. The state of the Primary and Secondary RLINKs is changed to STALE. The RVG state does not change, but the RVG is put into PASSTHRU mode that allows update of the Primary volume to continue until the error is fixed. See “[RVG PASSTHRU Mode](#)” on page 315.

The SRL must be repaired manually and then associated with the RVG. While the SRL is being repaired, no attempt is made to send data to the RLINKs. After the SRL is replaced, all RLINKs must be completely synchronized. Attach the RLINKs and perform a complete synchronization of the Secondaries.

On the Primary (seattle):

### ▼ To cleanup after a Primary SRL error

1. Dissociate the SRL from the RVG.

```
# vxvol -g hrdg dis hr_srl
```

2. Fix or replace the SRL volume.

3. Make sure that the repaired SRL is started before associating it with the RVG. If the repaired SRL is not started, start it:

```
# vxvol -g hrdg start hr_srl
```

4. Associate a new SRL with the RVG. After associating the new SRL, the RVG PASSTHRU mode no longer displays in the output of the command `vxprint -lv`.

```
# vxvol -g hrdg aslog hr_rvg hr_srl
```

5. Completely synchronize the Secondary. See “[Synchronizing the Secondary and Starting Replication](#)” on page 58 for details.

## RVG PASSTHRU Mode

Typically, writes to data volumes associated with an RVG go to the RVG's SRL first, and then to the RLINKs and data volumes. If the Primary SRL is ever detached because of an access error, then the Primary RVG is put into PASSTHRU mode. In PASSTHRU mode, writes to the data volume are passed directly to the underlying data volume, bypassing the SRL. No RLINKs receive the writes. Use `vxprint -l` on the RVG to see if the `passthru` flag is set. Associating a new SRL will clear PASSTHRU mode, and the Secondary node RVGs must be synchronized.

## Primary SRL Volume Error at Reboot

If the Primary SRL has an error during reboot, there is a possibility that the disks or arrays containing the SRL have not yet come online. Because of this, instead of placing the RVG in PASSTHRU mode, VVR does not recover the RVG. When the SRL becomes available, issue the following commands to recover the RVG and the RLINK:

```
# vxrvg -g diskgroup recover rvg_name
# vxrlink -g diskgroup recover rlink_name
```

After this error has occurred and you have successfully recovered the RVG, if you dissociate a volume from the RVG, you may see the following message:

```
Because there could be outstanding writes in the SRL, the data volume
being dissociated should be considered out-of-date and inconsistent
```

You can ignore this message.

If the SRL is permanently lost, create a new SRL as described in “[Recovering from SRL Header Error](#)” on page 317. In this case, it is possible that writes that had succeeded on the old SRL and acknowledged to the application, were not yet flushed to the data volumes and are now lost. Consequently, you must restore the data volumes from backup before proceeding. Because this causes the data volumes to be completely rewritten, it is recommended that you detach the RLINKs and synchronize them after the restore operation is complete.



## Primary SRL Volume Overflow Recovery

Because the size of the Primary SRL is finite, prolonged halts in update activity to any RLINK can exceed the log's ability to maintain all the necessary update history to bring an RLINK up-to-date. When this occurs, the RLINK in question is marked as STALE and requires manual recovery before replication can proceed. A STALE RLINK can only be brought up-to-date by using automatic synchronization or a block-level backup and checkpoint. The other RLINKs, the RVG, and the SRL volume are all still operational.

SRL overflow protection can be set up to prevent SRL overflow, and is the default. Instead of allowing the RLINK to become STALE, dcm logging is initiated. At a later time when the communication link is not overloaded, you can incrementally resynchronize the RLINK using the `vradm resync rvg` command.

## Primary SRL Header Error Cleanup and Recovery

An SRL header failure on the Primary is a serious error. All RLINKs are lost and must be recovered using a Primary checkpoint. Because information about data volume errors is kept in the SRL header, the correct status of data volumes cannot be guaranteed under all occurrences of this error. For this reason, we recommend that the SRL be mirrored.

If an SRL header error occurs during normal operation and you notice it before a reboot occurs, you can be certain that any data volumes that have also (simultaneously) failed will have a status of DETACHED. If the system is rebooted before the `vxprint` command shows the volumes to be in the DETACHED state, the status of any failed data volumes may be lost. Both these cases involve multiple errors and are unlikely, but it is important to understand that the state of Primary data volumes can be suspect with this type of error.

When a Primary SRL header error occurs, writes to the RVG continue; however, all RLINKs are put in the STALE state. The RVG is operating in PASSTHRU mode.

## Recovering from SRL Header Error

To recover from an SRL header error, dissociate the SRL from the RVG, repair the SRL, and completely synchronize all the RLINKs.

1. Stop the RVG.

```
# vxrvrg -g hrdg stop hr_rvg
```

2. Dissociate the SRL from the RVG.

```
# vxvol -g hrdg dis hr_srl
```

3. Repair or restore the SRL. Even if the problem can be fixed by repairing the underlying subdisks, the SRL must still be dissociated and reassociated to initialize the SRL header.

4. Make sure the SRL is started, and then reassociate the SRL:

```
# vxvol -g hrdg start hr_srl  
# vxvol -g hrdg aslog hr_rvg hr_srl
```

5. Start the RVG:

```
# vxrvrg -g hrdg start hr_rvg
```

6. Restore the data volumes from backup if needed. Synchronize all the RLINKs. See [“Methods to Synchronize the Secondary”](#) on page 58.

## Secondary Data Volume Error Cleanup and Recovery

If an I/O error occurs during access of a Secondary data volume, the data volume is automatically detached from the RVG and the RLINKs are disconnected. A subsequent attempt by the Primary to connect to the Secondary fails and a message that the Secondary volumes are stopped is displayed. The Primary is unaffected and writes continue to be logged into the SRL. After the Secondary data volume error is fixed and the data volume is started, the RLINKs automatically reconnect.

If there is no suitable Primary or Secondary checkpoint, detach the RLINKs on both the Primary and Secondary, and then synchronize the RLINKs. See [“Restoring the Secondary from Online Backup”](#) on page 192 for details.



## Recovery Using a Secondary Checkpoint

This section explains how to recover from a Secondary data volume error using a Secondary checkpoint.

On the Secondary (london):

1. Repair the failed data volume. You need not dissociate the data volume if the problem can be fixed by repairing the underlying subdisks.

2. Make sure that the data volume is started:

```
# vxvol -g hrdg start hr_dv01
```

3. Restore data from the Secondary checkpoint backup to all the volumes. If all volumes are restored from backup, the Secondary will remain consistent during the synchronization. Restore the RLINK by issuing the following command:

```
# vxrlink -g hrdg -c sec_chkpt restore rlk_seattle_hr_rvg
```

## Cleanup Using a Primary Checkpoint

On the Secondary (london):

1. Repair the failed data volume as above. Be sure that the data volume is started before proceeding:

```
# vxvol -g hrdg start hr_dv01
```

2. Detach the RLINK to enable writing to the Secondary data volumes:

```
# vxrlink -g hrdg det rlk_seattle_hr_rvg
```

3. Restore data from the Primary checkpoint backup to all data volumes. Unlike restoration from a Secondary checkpoint, the Primary checkpoint data must be loaded onto all Secondary data volumes, not just the failed volume. If a usable Primary checkpoint does not already exist, make a new checkpoint. For details, see [“Example—Synchronizing the Secondary Using Block-level Backup”](#) on page 67.

4. Reattach the RLINK.

```
# vxrlink -g hrdg att rlk_seattle_hr_rvg
```

On the Primary (seattle):

Detach the RLINK and then reattach it from the Primary checkpoint using the following commands:

```
# vxrlink -g hrdg det rlk_london_hr_rvg
# vxrlink -g hrdg -c primary_checkpoint att rlk_london_hr_rvg
```

## Secondary SRL Volume Error Cleanup and Recovery

The Secondary SRL is used only during atomic recovery of an RLINK and when an IBC is active. If I/O errors occur during recovery of the Secondary SRL, the recovery fails, the SRL volume is automatically detached, and the RLINK is forced to the PAUSE state. Manual intervention is required to repair the physical problem, reattach the SRL, and resume the RLINK. Upon resumption, an automatic recovery of the RVG is retried and if it succeeds, update activity can continue. The only problem occurs if the Primary SRL overflows before the repair is complete, in which case a full synchronization is required.

If an error occurs in the data portion of the SRL, the RLINK is forced to the PAUSE state with the `secondary_paused` flag set. The SRL is not dissociated.

If an error occurs in the SRL header, the Secondary RVG is forced to the FAIL state and the SRL is dissociated.

On the Secondary (london):

1. Dissociate the SRL, fix it, and then re-associate it. The dissociation and re-association is necessary even if the problem can be fixed by repairing the underlying subdisks because this sequence initializes the SRL header.

```
# vxvol -g hrdg dis hr_srl
```

Fix or replace the SRL. Be sure the SRL is started before associating it:

```
# vxvol -g hrdg start hr_srl
```

```
# vxvol -g hrdg aslog hr_rvg hr_srl
```

2. Run the RLINK resume operation to clear the `secondary_log_err` flag.

```
# vxrlink -g hrdg resume rlk_seattle_hr_rvg
```



## Secondary SRL Header Error Cleanup and Recovery

An SRL header failure on the Secondary puts the Secondary RVG into the FAIL state, and sets the RLINK state to the PAUSE state on both the Primary and Secondary. Because information about data volume errors is kept in the SRL header, the correct state of data volumes is not guaranteed in all cases. If a Secondary SRL header failure occurs during normal operation and is noticed before a reboot occurs, any data volumes that also failed will have a state of DETACHED. If the system is rebooted before the `vxprint` command shows the volumes to be in the DETACHED state, the status of any failed data volumes may be lost. Both these cases involve multiple errors and are unlikely, but it is important to understand that the state of Secondary data volumes can be suspect with this type of error.

1. Dissociate the SRL volume.

```
# vxvol -g hrdg dis hr_srl
```

2. Repair the SRL volume. Even if the problem can be fixed by repairing the underlying subdisks, the SRL volume must still be dissociated and re-associated to initialize the SRL header.

3. Start the SRL volume. Then, re-associate it.

```
# vxvol -g hrdg start hr_srl
# vxvol -g hrdg aslog hr_rvg hr_srl
```

4. Start the RVG.

```
# vxrvrg -g hrdg start hr_rvg
```

5. If the integrity of the data volumes is not suspect, just resume the RLINK.

```
# vxrlink -g hrdg resume rlk_seattle_hr_rvg
```

OR

If the integrity of the data volumes is suspect, and a Secondary checkpoint backup is available, restore from the Secondary checkpoint.

```
# vxrlink -g hrdg det rlk_seattle_hr_rvg
# vxrlink -g hrdg -f att rlk_seattle_hr_rvg
# vxrlink -g hrdg -w pause rlk_seattle_hr_rvg
```



---

Restore the Secondary checkpoint backup data on to the data volumes.

```
# vxrlink -g hrdg -c secondary_checkpoint restore  
rlk_seattle_hr_rvg
```

OR

If the integrity of the data volumes is suspect and no Secondary checkpoint is available, synchronize the Secondary using a block-level backup and Primary checkpoint (see [“Example—Synchronizing the Secondary Using Block-level Backup”](#) on page 67) or automatic synchronization.

```
# vxrlink -g hrdg det rlk_seattle_hr_rvg
```

On the Secondary, restore the Primary checkpoint backup data to the data volumes.

```
# vxrlink -g hrdg -f att rlk_seattle_hr_rvg
```

On the Primary (seattle):

```
# vxrlink -g hrdg -c primary_checkpoint att rlk_london_hr_rvg
```



## Secondary SRL Header Error at Reboot

If the secondary SRL has an error after a reboot, it is not possible to recover it, even if the SRL subsequently becomes available. Ignore the following message:

```
VxVM VVR vxrvlg ERROR V-5-1-0 RVG rvlg_name cannot be recovered
because SRL is not accessible. Try recovering the RVG after the SRL
becomes available using vxrecover -s command
```

1. Dissociate the SRL:

```
# vxvol -g hrdg -f dis srl
```

Ignore the following messages:

```
VxVM vxvol WARNING V-5-1-0 WARNING: Rvg rvlgname has not been
recovered because the SRL is not available. The data volumes may be
out-of-date and inconsistent
VxVM vxvol WARNING V-5-1-0 The data volumes in the rvg rvlgname
cannot be recovered because the SRL is being dissociated. Restore
the data volumes from backup before starting the applications
```

2. Create a new SRL volume, *new\_srl* and continue as follows:

```
# vxvol -g hrdg aslog rvlg_name new_srl
# vxrlink -g hrdg recover rlink_name
# vxrlink -g hrdg -f att rlink_name
# vxrvlg -g hrdg start rvlg_name
```

If replication was frozen due to receipt of an IBC, the data in the SRL is lost but there is no indication of this problem. To see whether this was the case, examine the `/var/adm/syslog/syslog.log` file for a message such as:

```
WARNING: VxVM VVR vxio V-5-0-0 Replication frozen for rlink <rlink>
```

If this is the last message for the RLINK, that is, if there is no subsequent message stating that replication was unfrozen, the Primary RLINK must be completely resynchronized.

# VVR Command Reference



This appendix lists the VVR commands and their descriptions. The `vradmin` command can be issued from any host in the Replicated Data Set (RDS); the low-level VVR commands must be issued on the host on which the object resides.

**Note** This reference lists command options for frequently used scenarios. For a complete list of options, refer to the respective manual page.

## Command Reference

VVR Command	Command Description
<code>vradmin -g diskgroup createpri rvg_name dv01_name,dv02_name... srl_name</code>	Creates Primary RVG of an RDS.
<code>vradmin -g diskgroup addsec local_rvgname pri_hostname sec_hostname</code>	Adds a Secondary RVG to an RDS.
<code>vradmin -g diskgroup set local_rvgname sec_hostname synchronous=value</code>	Sets up mode of replication: synchronous=off sets asynchronous synchronous=override sets synchronous vradmin command does not allow you to set synchronous=fail. You can do this using the vxedit command. For more information on the vxedit command refer to the vxedit manual page.



VVR Command	Command Description
<code>vradmin -g diskgroup set local_rvgname sec_hostname latencyprot=value</code>	Sets up Latency Protection: latencyprot=fail latencyprot=override latencyprot=off
<code>vradmin -g diskgroup set local_rvgname sec_hostname latency_high_mark=n</code>	Sets up latency_high_mark: latency_high_mark=n
<code>vradmin -g diskgroup set local_rvgname sec_hostname latency_low_mark=n</code>	Sets up latency_low_mark: latency_low_mark=n
<code>vradmin -g diskgroup set local_rvgname sec_hostname srlprot=value</code>	Sets up SRL Overflow Protection: srlprot=autodcm (default) srlprot=dcm srlprot=override
<code>vradmin -g diskgroup set local_rvgname sec_hostname packet_size=n</code>	Sets up the packet size.
<code>vradmin -g diskgroup set local_rvgname sec_hostname protocol=value</code>	Sets the protocol.
<code>vradmin -g diskgroup changeip local_rvgname [sec_hostname] attrs....</code>	Changes the host name or IP address of the Primary and Secondary RLINKs to the new values specified in the <i>newpri</i> and <i>newsec</i> attributes.
<code>vradmin -g diskgroup -l repstatus local_rvgname</code>	Displays consolidated replication-related information about an RDS.
<code>vradmin [-l] printrvg</code>	Displays information for all RDSs on local host.
<code>vradmin -g diskgroup [-l] printrvg local_rvgname</code>	Displays detailed information for a specific RDS.
<code>vradmin printvol</code>	Displays information about data volumes in all RDSs on the local host.
<code>vradmin -g diskgroup printvol local_rvgname</code>	Displays information about data volumes in an RDS.

VVR Command	Command Description
<code>vradmin -g diskgroup pauserep local_rvgname sec_hostname</code>	Pauses replication to a Secondary.
<code>vradmin -g diskgroup resumerep local_rvgname sec_hostname</code>	Resumes replication to a Secondary.
<code>vradmin -g diskgroup -a startrep local_rvgname sec_hostname</code>	Starts replication and synchronizes the Secondary using autosync.
<code>vradmin -g diskgroup -c checkpt_name startrep local_rvgname sec_hostname</code>	Starts replication and synchronizes the Secondary using a checkpoint.
<code>vradmin -g diskgroup stoprep local_rvgname sec_hostname</code>	Stops replication to a Secondary.
<code>vradmin -g diskgroup -c checkpt_name syncrvg local_rvgname sec_hostname....</code>	Synchronizes the Secondary volumes with the corresponding Primary volumes based on differences when the application is active or inactive.
<code>vradmin -g diskgroup -full -c checkpt_name syncrvg local_rvgname sec_hostname....</code>	Performs full synchronization with checkpoint to the Secondary when the application is active or inactive.
<code>vradmin -g diskgroup -full syncvol local_vols_list remote_hostname....</code>	Synchronizes volumes on local host and remote hosts using full synchronization.
<code>vradmin -g diskgroup -verify syncrvg local_rvgname sec_hostname...</code>	Verifies and reports any data differences between Secondary volumes and the corresponding Primary volumes.
<code>vradmin -g diskgroup -verify syncvol local_vols_list remote_hostname...</code>	Verifies and reports any data differences between remote volumes and the corresponding local volumes.



VVR Command	Command Description
<pre>vradmin -g diskgroup [-k {cache snap}] verifydata local_rvgname sec_hostname {cache=cacheobj   cachesize=size}</pre>	<p>Verifies that the data on the Secondary data volumes is identical to the Primary data volumes.</p> <p><b>Note</b> The <code>-k</code> option cannot be used if you are using the <code>cachesize</code> option to create a cache object as these cache objects are destroyed automatically once the <code>vradmin verifydata</code> command executes successfully.</p>
<pre>vradmin -g diskgroup addvol local_rvgname volume_name</pre>	<p>Adds a volume to an RDS.</p>
<pre>vradmin -g diskgroup [-f] resizevol local_rvgname volume_name length</pre>	<p>Resizes a data volume in an RDS.</p>
<pre>vradmin -g diskgroup delvol local_rvgname volume_name</pre>	<p>Removes a data volume from an RDS.</p>
<pre>vradmin -g diskgroup ibc local_rvgname task_name [sec_host]...</pre>	<p>Performs the specified off-host processing task on the Secondary.</p>
<pre>vradmin -g diskgroup migrate local_rvgname newprimary_name</pre>	<p>Migrates the Primary role to Secondary <code>newprimary_name</code>.</p>
<pre>vradmin -g diskgroup takeover local_rvgname</pre>	<p>Takes over the Primary role from an original Primary with fast failback enabled.</p>
<pre>vradmin -g diskgroup -autofb takeover local_rvgname</pre>	<p>Takes over the Primary role from an original Primary with fast failback enabled and automatically synchronizes the original Primary when it becomes available.</p>
<pre>vradmin -g diskgroup -N takeover local_rvgname</pre>	<p>Changes the role from Secondary to Primary without enabling fast failback.</p>

VVR Command	Command Description
<pre>vradmin -g diskgroup fbsync local_rvgname [cache=cache-object   cachesize=size]</pre>	<p>Converts the original Primary to a Secondary and starts resynchronization of the original Primary using fast-failback. Optionally, it also takes space-optimized snapshots of the original Primary's data volumes before starting the resynchronization.</p>
<pre>vradmin -g diskgroup -wait fbsync local_rvgname</pre>	<p>Converts the original Primary to a Secondary and starts resynchronization of the original Primary using fast failback. The command returns after resynchronization completes.</p>
<pre>vradmin -g diskgroup makesec local_rvgname newprimary_name</pre>	<p>Converts an original Primary to a Secondary when fast failback was not enabled.</p>
<pre>vradmin -g diskgroup resync local_rvgname [cache=cache-object   cachesize=size]</pre>	<p>Replays a DCM that is active due to an SRL overflow to incrementally synchronize the Secondary. Optionally, it also takes space-optimized snapshots of the original Primary's data volumes before starting the resynchronization.</p>
<pre>vradmin -g diskgroup -wait resync local_rvgname</pre>	<p>Replays a DCM that is active due to an SRL overflow to incrementally synchronize the Secondary. The command returns after resynchronization completes.</p>
<pre>vradmin -g diskgroup delsec local_rvgname sec_hostname</pre>	<p>Removes a Secondary from an RDS.</p>
<pre>vradmin -g diskgroup delpri rvg_name</pre>	<p>Removes a Primary when the application is inactive.</p>
<pre>vradmin -g diskgroup -f delpri rvg_name</pre>	<p>Removes a Primary when the application is active.</p>



VVR Command	Command Description
<code>vxrvg -g diskgroup [-1] getdatavols rvg_name</code>	Displays the names of all data volumes that are associated with the specified RVG.
<code>vxrvg -g diskgroup [-1] getrlinks rvg_name</code>	Displays the names of all RLINKs that are associated with the specified RVG.
<code>vxrvg -g diskgroup start rvg_name</code>	Enables I/O access to the data volumes associated with the RVG.
<code>vxrvg -g diskgroup stop rvg_name</code>	Disables I/O access to the data volumes associated with the RVG.
<code>vxrvg -g diskgroup recover rvg_name</code>	Recovers the RVG after rebooting a node.
<code>vxrvg -g diskgroup -c checkpt_name checkstart rvg_name</code>	Marks the beginning of a Primary checkpoint by pointing to the current location of the SRL.
<code>vxrvg -g diskgroup checkend rvg_name</code>	Marks the end of a Primary checkpoint by pointing to the current location of the SRL.
<code>vxrvg -c checkpt_name checkdelete rvg_name</code>	Deletes the specified Primary checkpoint.
<code>vxrvg -g diskgroup cplist rvg_name</code>	Displays information about all existing checkpoints that are associated with the RVG.





VVR Command	Command Description
<pre>vxrvg -g diskgroup [-f] [-p] [-P prefix   -a] snapback rvg_name</pre>	<p>Reattaches snapshot volumes to their original volumes in the RVG. This operation is similar to the <code>vxassist snapback</code> command for traditional (third-mirror breakoff) snapshots, and the <code>vxsnap reattach</code> command for instant snapshots.</p> <p><b>Note</b> When reattaching snapshots to an RVG that consists of only VxVM ISP volumes you must use the <code>-f</code> option otherwise the command will fail.</p>
<pre>vxrvg -g diskgroup snapprint rvg_name</pre>	<p>Displays information on the relationships between the data volumes of an RVG and any corresponding snapshots.</p>
<pre>vxrvg -g diskgroup [-P prefix] snaprefresh rvg_name</pre>	<p>Refreshes all existing snapshot volumes from the corresponding data volumes in the specified RVG. If a prefix is specified, this command refreshes only the snapshot volumes with that prefix. This creates a new point-in-time image for each snapshot, and the snapshot is immediately available for use with its new contents.</p>
<pre>vxrvg -g diskgroup [-f] [-P prefix] snaprestore rvg_name</pre>	<p>Restores the contents of all of the data volumes in the specified RVG from the corresponding snapshot volumes. If a prefix is specified, this command restores the contents only from the snapshot volumes with that prefix. The volumes are immediately available with the restored contents.</p>
<pre>vxrvg -g diskgroup [-i interval] [-t timestamp_frequency] [-C count] stats rvg_name</pre>	<p>Displays the application statistics for the specified RVG. This is only valid on the Primary.</p>



VVR Command	Command Description
<pre>vxrvg -g diskgroup [-P prefix] [-F -S] snapshot rvg_name [instantfull=volume_list] [instantso=volume_list] [plexbreakoff=volume_list] [exclude=volume_list] [plexprefix=plex_prefix] [cache=cachename  cachesize=size] [syncing={yes no}] [comment="comment"]</pre>	<p>Creates snapshots for all volumes in the specified RVG. This operation is similar to the vxassist snapshot command for traditional (third-mirror breakoff) snapshots, and the vxsnap make command for instant snapshots.</p> <p><b>Note</b> If the volumes in an RVG are VxVM ISP volumes then only instant full and instant space optimized snapshots are supported.</p>
<pre>vxrlink -g diskgroup assoc rvg_name rlink_name</pre>	<p>Associates an RLINK with an RVG.</p>
<pre>vxrlink -g diskgroup dis rlink_name</pre>	<p>Dissociates an RLINK from the RVG with which it is associated.</p>
<pre>vxrlink -g diskgroup [-a -c checkpoint_name]  -f] att rlink_name</pre>	<p>Enables an RLINK to connect to its remote RLINK by using auto attach, checkpoint attach or force attach.</p>
<pre>vxrlink -g diskgroup det rlink_name</pre>	<p>Detaches an RLINK.</p>
<pre>vxrlink -g diskgroup pause rlink_name</pre>	<p>Pauses updates to the Secondary RVG.</p>
<pre>vxrlink -g diskgroup resume rlink_name</pre>	<p>Resumes updates to the Secondary RVG that has been paused.</p>
<pre>vxrlink -g diskgroup recover rlink_name</pre>	<p>Recovers the RLINK after rebooting a node.</p>
<pre>vxrlink -g diskgroup -c checkpoint_name restore rlink_name</pre>	<p>Restores a failed Secondary RVG from a previously taken backup and Secondary checkpoint.</p>
<pre>vxrlink -c checkpoint_name checkdelete rlink_name</pre>	<p>Deletes the specified Secondary checkpoint.</p> <p><b>Note</b> This command must be run only on the Primary.</p>

VVR Command	Command Description
<code>vxrlink -g diskgroup verify rlink_name</code>	Displays the configuration status of the given RLINK.
<code>vxrlink -g diskgroup [-e] stats rlink_name</code>	Gives details of the use of the network by VVR. The -e option displays extended statistics.
<code>vxrlink -g diskgroup [-i &lt;interval&gt;] [-T] status rlink_name</code>	Displays how much of the SRL is being used by the RLINK and how much the Secondary is behind. This incremental synchronization status is displayed after an interval of <i>&lt;i&gt;</i> seconds. This command output can also display the status with a timestamp, using the -T option.
<code>vxrlink -g diskgroup cplist rlink_name</code>	Displays information about the existing Secondary checkpoints associated with the RLINK, which includes the name of the checkpoint, its size, and the percentage of SRL used.
<code>vxrlink -g diskgroup [-T] updates rlink_name</code>	Valid on the Secondary only. Displays the update ID received by the Secondary, as well as the number of updates by which the Primary is ahead. This information can be used to determine the most up-to-date Secondary RVG. This command when used with the -T option, displays the exact time in hours by which the Secondary is behind.
<code>vrstat -g diskgroup [-R] [-V] [-M] rvg_name</code>	Valid on the Primary as well as the Secondary. Displays detailed statistics. The output of this command is a combination of the outputs of the <code>vxrlink stats</code> , <code>vxrlink status</code> , <code>vxstat</code> , and <code>vxmemstat</code> commands.
<code>vxprint -V[1]</code>	Displays all RVGs.



VVR Command	Command Description
<code>vxprint -P[1]</code>	Displays all RLINKS.
<code>vxmake -g diskgroup rlink rlink_name protocol=protocol_name remote_host=sec_hostname remote_rlink=rlink_name</code>	Creates an RLINK with the specified network transport protocol. The attribute <i>protocol_name</i> can have a value of <i>TCP</i> or <i>UDP</i> .
<code>vxmemstat [-i interval [-t count]] [-e]</code>	Display memory statistics for VERITAS Volume Manager.
<code>vxtune [ -rH ] keyword arg ...</code>	Modify and display Volume Replicator and VxVM tunables.
<code>vrport [ -a   -r ] keyword arg ...</code>	Perform Volume Replicator port management operations.
<code>vrnotify -g diskgroup [-n number] [-t timeout] [rvg_name ...]</code>	Display VERITAS Volume Replicator (VVR) events.
<code>vxedit -g diskgroup set protocol=protocol_name rlink_name</code>	Allows you to change the specified network transport protocol. The protocol can be set to either <i>TCP</i> or <i>UDP</i> . For more information refer to the <code>vxedit</code> manual page.

# Messages



This appendix contains a listing of diagnostic messages generated by VERITAS Volume Replicator. The messages are classified into Kernel and Utility messages. The Kernel messages are categorized by type, for example, RLINK, SRL and DCM, and communication errors.

## Kernel Messages

This section describes the messages displayed by the VVR Kernel that appear on the console and in the `/var/adm/syslog/syslog.log` file. Error messages are listed first, followed by a list of informational messages.

## Error Messages

This section lists messages related to RLINKs, the SRL and the DCM, communication errors, configuration errors, I/O failure, and kernel logging.

### Messages Related to RLINKs

Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-114	Disconnecting rlink <i>rlink_name</i> to permit transaction to proceed.	Ignore message if seen occasionally. <b>Action:</b> If the message persists, contact VERITAS Customer Support.



Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-208	Log over 80% full for rlink <i>rlink_name</i>	This indicates that SRL is almost full. <b>Action:</b> Check that the <code>srlprot</code> attribute is set correctly in order to avoid failing of writes, throttling of writes or overflow of SRL.
V-5-0-267	Rlink <i>rlink</i> disconnecting due to ack timeout on <i>msg type</i> message	Ignore message if seen occasionally. <b>Action:</b> If the message persists, contact VERITAS Customer Support.
V-5-0-329	Unable to connect rlink <i>rlink_name</i> on <i>rvg rvg_name</i> : Unknown error ( <i>errno</i> )	VVR encountered an unknown error. <b>Action:</b> Contact VERITAS Customer Support.
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on <i>rvg rvg_name</i> : Disk group or rlink not found on remote	The Primary RLINK indicates a Secondary disk group and RLINK combination that does not exist on the Secondary. <b>Action:</b> Ensure that the disk group and the RLINK specified by the Primary RLINK exist on the Secondary.
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on <i>rvg rvg_name</i> : Rlink detached on remote	The RLINK on the Secondary node is in the DETACHED or STALE state. <b>Action:</b> Attach the RLINK and retry.



Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : Not ready on remote	Ignore message if seen occasionally. <b>Action:</b> If the message persists, contact VERITAS Customer Support.
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : Rlink already connected on remote	
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : Stream error on remote	
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : Checksum error on remote	
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : Unexpected command on remote	
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : Out of space on remote	
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : Port closing on remote	
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : Too many threads on remote	
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : Invalid port on remote	
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : Send error on remote	



Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-330	Unable to connect rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : KTLI connect failed	Ignore message if seen occasionally. <b>Action:</b> If the message persists, contact VERITAS Customer Support.
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : Time out on remote	The Secondary machine is unreachable. Will clear up when the network or the Secondary node recovers.
V-5-0-0	Disconnecting rlink <i>rlink</i> due to error in sending ( <i>error-code</i> )	Ignore message if seen occasionally. <b>Action:</b> If the message persists, contact VERITAS Customer Support.
V-5-0-0	Disconnecting rlink <i>rlink</i> due to loss of TCP connection	
V-5-0-0	Disconnecting rlink <i>rlink</i> due to stream error	
V-5-0-0	Disconnecting rlink <i>rlink</i> due to header checksum error	
V-5-0-0	Disconnecting rlink <i>rlink</i> due to unexpected message	
V-5-0-0	Disconnecting rlink <i>rlink</i> as Secondary data volumes are stopped	
V-5-0-0	Disconnecting rlink <i>rlink</i> due to bad message.	
V-5-0-0	Disconnecting rlink <i>rlink</i> due to error : <i>error-code</i>	
V-5-0-0	Disconnecting rlink <i>rlink</i> remote already connected	





Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-0	Disconnecting rlink <i>rlink</i> due to pending transaction	Ignore message if seen occasionally. <b>Action:</b> If the message persists, contact VERITAS Customer Support.
V-5-0-0	Received from unexpected port. Expected from port <i>port</i> , received from port <i>port</i>	
V-5-0-0	Header checksum error	
V-5-0-0	Received unexpected message. Expected message with opcode <i>operation-code</i> , received opcode <i>operation-code</i>	
V-5-0-0	Data checksum error for handshake message id ( <i>message-id</i> ), data checksum computed ( <i>checksum-value</i> ) but header contains ( <i>checksum-value</i> )	
V-5-0-0	Data checksum error. Received message id <i>message-id</i> with data checksum : <i>checksum-value</i> expected checksum : <i>checksum-value</i>	



## Messages Related to the SRL and DCM

Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-100	DCM Logs not accessible, dcm logging aborted	DCM logs are not accessible. This might have happened due to a media failure, in which case, other errors may have been logged to the console. This error is unlikely to occur, because the default is to mirror the DCM.
V-5-0-102	DCM volume <i>vol name</i> is detached	The DCM volume became detached because a DCM log entry cannot be written. This might have happened due to a media failure, in which case other errors may have been logged to the console. Any RLINKs currently using the DCM will be detached and marked STALE. This error is unlikely to occur, because the default is to mirror the DCM.
V-5-0-107	Detaching rlink <i>rlink_name</i> due to I/O error on remote SRL during recovery	When the original Primary comes up after it has been taken over, an I/O error has occurred on the original Primary's SRL while performing recovery. The new Primary has detached the RLINK that connects to the failed Primary because now there is no way to make the old Primary consistent and convert it to a Secondary of the new Primary. <b>Action:</b> Fix the I/O error on the SRL, convert the original Primary into a Secondary of the new Primary, and perform a complete synchronization.



Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-280	Rlink <i>rlink</i> STALE due to log overflow	The specified RLINK has fallen too far behind, and overflowed the SRL. It will be detached and marked STALE. The Secondary will require a complete resynchronization. This can be avoided in the future by using the RLINK's <code>srlprot</code> attribute.
V-5-0-287	<i>rvg rvg_name</i> , SRL <i>srl</i> : Inconsistent log - detaching all rlinks	SRL contains corrupt data, and so is not usable. All RLINKs are detached and marked STALE. All secondaries will require a complete resynchronization. This error probably indicates a bug in VVR. <b>Action:</b> Contact VERITAS Customer Support.
V-5-0-288	Secondary log overflowed. Pausing rlink <i>rlink_name</i>	Specified Secondary RLINK gets paused. <b>Action:</b> Associate the SRL of the same size as that of Primary to the Secondary RVG and resume the RLINK.
V-5-0-293	SRL for RVG <i>rvg_name</i> contains old version of SRL header	Specified SRL associated with the RVG is of older version. You may not have followed the proper procedure for upgrading VVR as described in Release Notes. <b>Action:</b> Refer to the latest Release Notes for the correct upgrading procedures.



<b>Unique Message Identifier (UMI)</b>	<b>Message</b>	<b>Message Definition</b>
V-5-1-435	Cannot allocate space for 20480 block volume or not enough disks for creating dcm with 2 mirrors	<p>By default, the DCM is mirrored; therefore, it requires space on two disks. An attempt to associate a DCM to a new data volume or an existing data volume failed because sufficient space is not available.</p> <p><b>Action:</b> We recommend that you provide the required space. Alternatively, you can specify <code>nlog=1</code> in the <code>vxassist make</code> or <code>vxassist addlog</code> command.</p>

## Messages Related to Communication Errors

Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-18	startdaemon_deferred: Could not create volkmsgd	Ignore message if seen occasionally. <b>Action:</b> If the message persists, contact VERITAS Customer Support.
V-5-0-44	Cannot alloc bind structure ( <i>errno</i> )	
V-5-0-45	Cannot alloc bind structure for listen server ( <i>errno</i> )	
V-5-0-46	Cannot alloc connection indication call structure ( <i>errno</i> )	
V-5-0-47	Cannot alloc ktli header ( <i>errno</i> )	
V-5-0-40	Cannot bind the acceptor ( <i>errno</i> )	
V-5-0-49	Cannot connect rlink <i>rlink_name</i> due to protocol mismatch: remote rlink is using <i>protocol_name</i> protocol, local rlink is using <i>protocol_name</i> protocol	
V-5-0-50	Cannot connect to host <i>ip address</i> port <i>port_number</i> ( <i>errno</i> )	You may see this message when <code>vxnetd</code> has not been run on the remote host. <b>Action:</b> Run <code>vxnetd</code> by issuing the command <code>/sbin/init.d/vxnm-vxnetd</code> . If the message persists, contact VERITAS Customer Support.
V-5-0-51	Cannot connect to remote due to header format mismatch or checksum error	<b>Action:</b> Contact VERITAS Customer Support.



Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-54	Cannot find any free port to bind	All the ports in the specified list have been utilized. <b>Action:</b> Add the ports to be used for replication using the <code>vrport</code> command.
V-5-0-75	Cannot open ktli ( <i>errno</i> )	Ignore message if seen occasionally. <b>Action:</b> If the message persists, contact VERITAS Customer Support.
V-5-0-76	Cannot open ktli for listen server ( <i>errno</i> )	
V-5-0-77	Cannot open the acceptor port ( <i>errno</i> )	
V-5-0-78	Cannot open the server port ( <i>errno</i> )	
V-5-0-80	Cannot send out a message	
V-5-0-82	Cannot spawn netiod thread; invalid transport selected	
V-5-0-84	Cannot spawn server thread	
V-5-0-611	Could not send heartbeat <i>id</i> to node <i>node</i>	
V-5-0-174	header checksum error	
V-5-0-175	Heartbeat unacknowledged from node <i>ip-addr</i> for <i>t</i> seconds	
V-5-0-206	listen server port in use ( <i>errno</i> )	Port assigned to VVR is in use by another application. <b>Action:</b> Change the assigned port using the <code>vrport</code> command.
V-5-0-219	nmcom client cannot bind ktli port ( <i>errno</i> )	Ignore message if seen occasionally. <b>Action:</b> If the message persists, contact VERITAS Customer Support.
V-5-0-220	nmcom client cannot open ktli port ( <i>errno</i> )	



Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-236	Number of ports available ( <i>total-port</i> ) is less than the number of rlinks ( <i>numreplicas</i> ) in the system	<p>The specified number of ports are less than total number of RLINKs in the system. Some of the RLINKs will be disconnected.</p> <p><b>Action:</b> For all RLINKs that are to take part in replication, configure at least <i>n</i> number of ports using the <code>vrport</code> command, where <i>n</i> is <math>\geq \text{num\_of\_rlink}</math> in the system.</p>
V-5-0-246	port <i>port-id</i> is in use by another application	<p>Port assigned to VVR is in use by another application.</p> <p><b>Action:</b> Change the assigned port using the <code>vrport</code> command.</p>
V-5-0-247	port in use ( <i>port id</i> )	<p>Port assigned to VVR is in use by another application.</p> <p><b>Action:</b> Change the assigned port using the <code>vrport</code> command.</p>
V-5-0-253	Received 100 duplicate packets. Check network configuration	<p>Indicates a possible network misconfiguration, such as multiple NICs with the same address.</p>
V-5-0-254	Received from unexpected port	
V-5-0-255	Received unexpected message	
V-5-0-376	VVR listener thread exiting	<p>If VVR is using the TCP protocol, a thread listens for incoming RLINK connections from remote hosts. When this thread encounters a fatal error, it exits with this message.</p> <p><b>Action:</b> Run <code>vxnetd</code> by issuing the command <code>/sbin/init.d/vxnm-vxnetd</code>. If the message persists, contact VERITAS Customer Support.</p>



Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-0	2a: Received invalid message block in TCP stream	Ignore message if seen occasionally. <b>Action:</b> If the message persists, contact VERITAS Customer Support.
V-5-0-0	2b: Received invalid message block in TCP stream	
V-5-0-0	2c: Received invalid message block in TCP stream	
V-5-0-0	2d: Received invalid message block in TCP stream	
V-5-0-0	Could not save message block, TCP stream error	
V-5-3-0	Received invalid message block in TCP stream	



## Messages Related to Configuration Errors

Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-29	Configuration error on rvg <i>rvg_name</i> pausing rlink <i>rlink_name</i>	Indicates configuration error. <b>Action:</b> Clear the mentioned error.
V-5-0-30	Cannot find matching volume for volume <i>pri_datavol</i> on primary	
V-5-0-31	Name on secondary <i>datavol_name</i> does not match name on primary <i>pri_datavol</i>	
V-5-0-32	Size of volume <i>pri_datavol</i> on primary ( <i>vol_size</i> ) does not match size on secondary ( <i>sec_vol_sz</i> )	
V-5-0-511	Replica has an unsupported message format	Remote host is running older version of VVR. The RLINKs will go into the STALE state. <b>Action:</b> For replication to take place, all the hosts in the VVR configuration must be running same version of VVR. The only exception being when Primary is running VVR 4.0 and the Secondary is running VVR 3.5 MP2.
V-5-0-512	Replica is running an unsupported version	
V-5-1-10128	Operation requires transaction	Indicates that another configuration change is in progress. <b>Action:</b> Issue the command again.



## Message Related to I/O Failure

Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-167	Failing writes on rvg <i>rvg_name</i> . Remote is now a primary rvg.	When the original Primary comes up after it has been taken over and the application on the original Primary tries to perform writes, the writes will fail because this RVG is no longer the Primary. <b>Action:</b> Stop the application on the original Primary.

## Messages Related to Kernel Logging

Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-59	cannot log <i>error</i> for rvg <i>rvg_name</i>	<p>The messages in this set mean that VVR attempted to log the specified event in the kernel log; however, the attempt failed. The attempted write to the log failed either because the kernel log is full or because of a write error to the drive.</p> <p><b>Action:</b> If error messages were seen from the disk driver, it is likely that the last copy of the log failed due to a disk error. The failed drive in the disk group must be replaced and the log will then be re-initialized on the new drive.</p> <p>If there were no error messages from the disk driver, contact VERITAS Customer Support.</p>
V-5-0-62	cannot log atomic commit changes for rvg <i>rvg_name</i>	
V-5-0-63	cannot log changes for rvg <i>rvg_name</i>	
V-5-0-65	cannot log datavol error for rvg <i>rvg_name</i>	
V-5-0-68	cannot log error during failback for rvg <i>rvg_name</i>	
V-5-0-70	cannot log srl error for rvg <i>rvg_name</i>	
V-5-0-71	cannot log srl error for rvg <i>rvg_name</i>	
V-5-0-72	cannot log unfreeze error for rvg <i>rvg_name</i>	
V-5-0-73	cannot log update commit state for rvg <i>rvg_name</i>	
V-5-0-95	could not log config error for rvg <i>rvg_name</i>	
V-5-0-160	failed to flush log on detach of rvg <i>rvg_name</i>	<p>The messages in this set mean that VVR logged the event(s) specified in the kernel log but was unable to flush the log.</p> <p><b>Action:</b> Contact VERITAS Customer Support.</p>
V-5-0-161	failed to flush log on errors on rvg <i>rvg_name</i>	



Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-165	Failed to log the detach of the DCM volume <i>vol name</i>	The messages in this set mean that VVR attempted to log the specified event in the kernel log; however, the attempt failed. The attempted write to the log failed either because the kernel log is full or because of a write error to the drive.  <b>Action:</b> If error messages were seen from the disk driver, it is likely that the last copy of the log failed due to a disk error. The failed drive in the disk group must be replaced and the log will then be re-initialized on the new drive.  If there were no error messages from the disk driver, contact VERITAS Customer Support.
V-5-0-348	vol_rp_set_state_atomic_commit_done:cannot log atomic commit state for rvg <i>rvg_name</i>	
V-5-0-349	vol_rp_set_state_atomic_commit_start : cannot log atomic commit state for rvg <i>rvg_names</i>	

## Informational Messages

The messages that appear in this section do not require an action. For the most part, they are informational.

### Informational RLINK Messages

Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-113	Disconnecting rep <i>rlink_name</i> to shift to DCM protection	This indicates a temporary disconnect during shifting to DCM protection.
V-5-0-149	Encountered IBC while flushing SRL to DCM - IBC dropped	The IBC residing on the SRL will be dropped when DCM protection gets triggered. You need to reissue the IBC after the DCM replay completes.
V-5-0-265	Rlink <i>rlink</i> connected to remote	The named RLINK has succeeded in connecting to its remote RLINK.
V-5-0-266	Rlink <i>rlink</i> disconnected from remote	<p>The named RLINK has disconnected from its remote RLINK. This can have many different causes, including:</p> <ul style="list-style-type: none"> <li>◆ The network is down</li> <li>◆ The Secondary node is down</li> <li>◆ The Primary RLINK was paused or detached</li> <li>◆ The Secondary RLINK was detached</li> </ul> <p>The RLINK will connect automatically when the situation clears.</p>



V-5-0-270	Rlink <i>rlink</i> has a secondary config error	Indicates different states of the specified RLINK.
V-5-0-271	Rlink <i>rlink</i> has a secondary log error	
V-5-0-274	Rlink <i>rlink</i> is in failed state	
V-5-0-275	Rlink <i>rlink</i> is inconsistent	
V-5-0-276	Rlink <i>rlink</i> is primary paused	
V-5-0-277	Rlink <i>rlink</i> is secondary paused	
V-5-0-278	Rlink <i>rlink</i> is stale and not replicating	
V-5-0-330	Unable to connect to rlink <i>rlink_name</i> on rvg <i>rvg_name</i> : Time out on remote	The Secondary machine is unreachable. Will clear up when the network or the Secondary node recovers.
V-5-0-467	Paused replay on RVG <i>rvg_name</i>	<p>The DCM replay gets paused on Primary. This can have many different causes, including:</p> <ul style="list-style-type: none"> <li>◆ The network is down</li> <li>◆ The Secondary node is down</li> <li>◆ The Primary RLINK was paused or detached.</li> <li>◆ The Secondary RLINK was paused or detached.</li> </ul> <p>DCM replay will start automatically when the RLINK gets connected.</p>
V-5-0-0	Disconnecting rlink <i>rlink</i> due to kernel reset	This may happen due to a user initiated action, which causes the kernel objects to be recreated.
V-5-0-0	Disconnecting rlink <i>rlink</i> to permit transaction to proceed	The RLINK may get temporarily disconnected if a transaction is in progress. Once the transactions is completed, the RLINK will get reconnected.

## Informational SRL and DCM Messages

Unique Message Identifier (UMI)	Message	Message Definition
V-5-0-79	cannot recover vol <i>vol_name</i>	During a reboot, the specified data volume was found to be disabled, and so cannot be recovered properly. The data volume may no longer be considered valid.
V-5-0-101	DCM replay complete on rlink <i>rlink_name</i>	DCM replay gets completed on the specified RLINK.
V-5-3-0	Resumed replay on RVG <i>rvg_name</i>	DCM replay gets resumed on the Primary. DCM replay will start automatically when the RLINK gets connected.



## Utility Error Messages

Utility messages are issued by the `vxconfigd` configuration daemon and are displayed on the terminal where the command was executed. The messages listed here are not all-inclusive. In most cases, the messages are self-explanatory. Only those messages that require a message definition are listed here.

Message	Message Definition
Out of kernel memory	The user request failed because the kernel could not allocate sufficient memory to complete the request successfully. If possible, attempt the request again. If it continues to fail, this may indicate that the system has insufficient resources.
IBC error	On the Secondary, an attempt was made to dissociate or delete an RLINK from an RVG against which there are outstanding IBC receive ioctls. All such outstanding ioctls must complete before the request will be processed.
Duplicate master_datavol mapping	Two volumes on the Secondary are mapped to the same volume on the Primary.
Multiple attached rlinks on secondary	There could be some error during the recovery. The Secondary can only have one active RLINK. Clear the utility fields and then dissociate any inactive RLINKs.
Replicated volume may not have a drl	An attempt was made to associate a volume with the DRL log to an RVG. A volume with a DRL log is not allowed to be associated to an RVG.
SRL cannot be resized	An attempt was made to resize the SRL. The resize operation is not allowed if the volume is an SRL.
V-5-1-3265 WARNING: DCM log size is smaller than recommended due to increased volume size.	To resize the DCM, remove it and use the <code>vxassist addlog</code> command to set the correct size.





<b>Message</b>	<b>Message Definition</b>
V-5-1-3265 WARNING: DCM log size will be smaller than recommended due to increased volume size.	Remove DCMs from volume and add them back specifying no <code>loglen</code> attribute, to get the default (recommended) size. Then re-issue the <code>vxassist maxgrow</code> command to find out how large the volume can be grown with the new DCM size.



## The vradmin Error Messages

This section lists and describes some of the error messages that may be displayed by `vradmin`, if a problem occurs. It also suggests the appropriate action that needs to be taken if a problem occurs.

Unique Message Identifier (UMI)	Message	Message Definition
V-5-36-2086	Server volume access error during [assign volids] volume path: [ <code>/dev/vx/dsk/dgname/volume name</code> ]	This error may occur while performing the <code>vradmin syncvol</code> command if the remote volume, that is, the volume being synchronized, is not enabled or available for read and write.  <b>Action:</b> Make sure the remote volume is available for read and write. Start the volume if it is currently stopped. If the volume is part of an RVG, make sure the RVG or the replication (i.e. RLINK) is in a certain state which allows read/write to the data volume.
V-5-52-12	<code>vradmind</code> server not running on this system.	<b>Action:</b> Start the <code>vradmind</code> server by running: <pre>/sbin/init.d/vras-vradmind.sh start</pre> Then, run the command again.
V-5-52-16	<code>vradmind</code> stopped running - Exiting.	The <code>vradmind</code> server got terminated during command execution. This could be due to an administrative action or due to some problem in <code>vradmind</code> .  <b>Action:</b> Contact VERITAS Customer Support to look into the problem. In the mean time, restart <code>vradmind</code> to resume normal operations.

V-5-84-162	Terminating the collection of statistics.	<p>This error may occur while trying to run the <code>vrstat</code> command on the secondary, if the <code>vradmind</code> server on the Primary was terminated while the <code>vrstat</code> command was in progress. This could be due to an administrative action or due to some problem in <code>vradmind</code>.</p> <p><b>Action:</b></p> <p>If you had stopped the <code>vradmind</code> server manually, then restart it and run the <code>vrstat</code> command again. If the <code>vradmind</code> server stopped due to some problem, contact VERITAS Customer Support.</p>
V-5-52-242	Attribute <code>cachesize</code> is not allowed with the <code>-k</code> cache option.	<p>When cache objects are created using the <code>cachesize</code> option they get deleted when snapshots are deleted.</p> <p><b>Action:</b></p> <p>Either use the <code>-k snap</code> option or create a cache object and use the <code>cache</code> attribute.</p>
V-5-52-401	RDS <code>rds</code> is processing another command or a configuration change.	<p>An attempt to run a <code>vradmind</code> command on RDS <code>rds</code> failed because another <code>vradmind</code> command is already running on this RDS or the RDS is undergoing a configuration change.</p> <p><b>Action:</b> Retry running the command that you had specified after some time.</p>
V-5-52-405	Primary <code>vradmind</code> server disconnected.	<p>The <code>vradmind</code> server on Primary is not running.</p> <p><b>Action:</b> Start the <code>vradmind</code> server on Primary.</p>
V-5-52-418	Volumes on host <code>host</code> are incorrectly mapped.	<p>One or more data volumes on the Primary do not have corresponding Secondary data volumes mapped to it.</p> <p><b>Action:</b> Check the volume mapping between the Primary and the corresponding Secondary using <code>vradmind printvol</code> command. There must be one-to-one mapping between Primary and Secondary data volumes.</p>



V-5-52-421	vradmind server on host <i>host</i> not responding or hostname cannot be resolved.	<p>The vradmind server on the host <i>host</i> is not running, or it's running on a different port, or <i>host</i> is not reachable.</p> <p><b>Action:</b> Make sure <i>host</i> is the correct host name and is reachable. Also make sure that vradmind is running on <i>host</i> and is using the same port as that on the Primary. Use <code>vrport</code> command to check and/or set the ports.</p>
V-5-52-422	Primary disk group <i>dg</i> is not authenticated.	<p>Before adding a new Secondary (using <code>addsec</code>) or overwriting data on a remote host (using <code>syncvol</code>), vradmind performs some authentication. This is done by checking that the <code>/etc/vx/vras/.rdg</code> file on the remote host contains an entry for the Primary disk group ID. The vradmind <code>addsec</code> or <code>syncvol</code> command fails if <code>/etc/vx/vras/.rdg</code> file on the remote host does not have such an entry.</p> <p><b>Action:</b> Add the Primary disk group ID to the <code>/etc/vx/vras/.rdg</code> file on the remote host. To find disk group ID, run the <code><b>vxprint -l diskgroup_name</b></code> command on the Primary.</p>



V-5-52-429	Secondary is processing a configuration change or its Primary is still reachable.	<p>This error may occur only when performing the <code>vradmin takeover</code> command when the Primary is still reachable from the Secondary or when the Secondary is undergoing a configuration change. To know if the Primary is reachable or not from this Secondary, run <code>vradmin -l printrvg rvg_name</code>. If the output of the <code>printrvg</code> command shows the <code>Pri</code> or <code>Sec</code> IP not available or <code>vradmind</code> not running or the <code>No Primary RVG</code> error in the Config Errors, it indicates that the Primary is not reachable from this Secondary.</p> <p><b>Action:</b> Check whether the Primary is reachable. If it is reachable then takeover cannot take place. In such a case, run the <code>vradmin migrate</code> command instead, to migrate the role of the Primary. If the Secondary is undergoing a configuration change, retry running the command after some time.</p>
V-5-52-447	RDS has configuration error. Check information about this RDS.	<p>Some <code>vradmin</code> commands will not proceed if the RDS has any configuration error(s).</p> <p><b>Action:</b> Check the Config Errors listing for the specific RDS in the output of the <code>vradmin -l printrvg</code> command to identify the errors. Correct the errors and then retry running the command.</p>
V-5-52-448	Secondary on <code>sec_host</code> has configuration error. Check information about this RDS.	<p>Some <code>vradmin</code> commands will not proceed if the specified Secondary has any configuration error(s). (The command will still proceed if any of the other Secondary RVGs in the RDS has a configuration error).</p> <p><b>Action:</b> Check the Config Errors listing for the specific RDS in the output of the <code>vradmin -l printrvg</code> command to identify the error. Correct the error and then retry running the command.</p>



V-5-52-449	Secondary <i>rvg_name</i> does not have an active Primary.	<p>The Secondary cannot determine its Primary due to one or more of the following reasons:</p> <ul style="list-style-type: none"> <li>◆ it has some configuration error.</li> <li>◆ the Primary is not reachable.</li> <li>◆ it does not have any RLINK.</li> </ul> <p><b>Action:</b> Run the <code>vradmin -l printrvg</code> on this Secondary to verify the cause of the problem.</p> <p>Correct the configuration error reported in the output of the <code>printrvg</code> command and then retry running the command.</p>
V-5-52-451	Volumes <i>vol_name</i> on Primary and Secondary differ in size; +/- not allowed in the size specification.	<p>When Primary and Secondary data volume <i>vol_name</i> differ in size then specifying a relative size using +/- is not allowed in the <code>vradmin resizevol</code> command.</p> <p>Use the <code>vradmin printvol</code> command to verify the size of the Primary and Secondary data volume.</p> <p><b>Action:</b> Specify absolute size for the volume instead of a relative size in the <code>vradmin resizevol</code> command.</p>
V-5-52-456	Volumes in RVG <i>rvg_name</i> do not have DCMs required for failback logging.	<p>This error occurs if the <code>vradmin takeover</code> command is run with failback logging enabled but at least one of the data volumes in the Secondary RVG does not have DCM. Failback logging requires that all volumes must have DCMs.</p> <p><b>Action:</b> If you want to run the <code>takeover</code> command with failback logging enabled, first add DCMs to all the data volumes. Then, run the <code>takeover</code> command. If you do not want to have failback logging enabled in the <code>takeover</code> command, then use it with the <code>-N</code> option. This requires either full synchronization or difference-based synchronization of the data volumes on the original Primary.</p>



V-5-52-465	Primary RVG <i>rvg_name</i> still has one or more Secondary RVGs.	<p>The <code>vradmin delpri</code> command performs the required deletion only when Primary RVG does not have any configured Secondaries.</p> <p><b>Action:</b> Before deleting a Primary RVG, you must delete all the Secondaries. Use the <code>vradmin delsec</code> command to delete a Secondary.</p>
V-5-52-467	Script <i>script_name</i> does not exist.	<p>If the missing script is the <code>onfreeze</code> script, then it must exist on the Secondary host, where the off-host processing task is performed.</p> <p><b>Action:</b> Put the <code>onfreeze</code> script on the Secondary. Make sure that the script is located in the following directory:  <code>/etc/vx/vvr/ibc_scripts/task_name</code>          Retry running the <code>vradmin ibc</code> command.</p>
V-5-52-468	<i>script_name</i> is not executable.	<p><b>Action:</b> Change the permission for the script <i>script_name</i> to make it executable by root user.</p>
V-5-52-469	Script <i>script_1</i> is provided; script <i>script_2</i> must also be provided.	<p>The <code>vradmin ibc</code> command requires that both the <code>quiesce</code> and the <code>unquiesce</code> script, if provided, must be present together on the Primary. If you only need one script, you will still need to provide the other one so that the <code>vradmin ibc</code> command executes properly. Therefore, for the script that you don't need create an empty script, that is, a script which contains only the following lines:</p> <pre>#!/sbin/sh exit 0</pre> <p><b>Action:</b> Create <i>script_2</i> on the Primary and retry running the <code>ibc</code> command.</p>



V-5-52-471	<i>host</i> has failback flag set.	Failback logging was enabled when the Secondary on <i>host</i> took over as a Primary. When the failback logging is enabled you cannot use the <code>makesec</code> command to convert original Primary to Secondary. <b>Action:</b> Use the <code>vradmin fbsync</code> command to convert the original Primary to a Secondary.
V-5-52-478	Cannot perform incremental synchronization: RLINK <i>rlk_name</i> on host <i>host</i> not in CONNECT state.	The <code>vradmin fbsync</code> command will proceed only if the RLINK <i>rlk_name</i> on the original Primary <i>host</i> is attached. This error occurs when the RLINK <i>rlk_name</i> somehow gets paused or is not in connect state. <b>Action:</b> Check the RLINK <i>rlk_name</i> on the original Primary <i>host</i> . If the RLINK is paused, resume the RLINK then retry running the <code>vradmin fbsync</code> command. If it is already detached, then failback synchronization cannot be performed. In this case you will need to use the <code>vradmin makesec</code> command to convert the original Primary to Secondary of the new Primary and then synchronize the original Primary (new Secondary) with the new Primary.
V-5-52-479	Host <i>host</i> not reachable.	This error generally occurs in either of the following: <ul style="list-style-type: none"> <li>◆ network connection to host <i>host</i> is down.</li> <li>◆ <code>vradmind</code> server on <i>host</i> is not running or is running on a different port.</li> </ul> <b>Action:</b> Correct the network if it's a network connection problem. In the case the <code>vradmind</code> server is not running start <code>vradmind</code> on host <i>host</i> using <pre>/sbin/init.d/vras-vradmind.sh start</pre> OR Make sure <code>vradmind</code> on <i>host</i> runs on the same port as the local <code>vradmind</code> . Use <code>vrport</code> command to check and reset the port.





V-5-52-480	Operation requires a disk group.	The command for which the specified operation failed requires a valid disk group to be present.
V-5-52-481	Cannot perform the operation in a cross-version replication environment.	In a cross-version replication environment, operations which can result in configuration changes, like <code>vradmin addvol</code> , <code>delvol</code> and <code>resizevol</code> , are not allowed. <b>Action:</b> Upgrade VVR on all the hosts running earlier version of VVR to the same version of VVR running on the other hosts in the RDS.
V-5-52-482	Cannot perform the operation: Secondary has an earlier version of VVR.	This error may occur in a cross-version replication environment while: <ul style="list-style-type: none"> <li>◆ Performing the <code>vradmin addsec</code> command to add a host running an earlier version of VVR to an RDS which has at least one host running a later version of VVR.</li> <li>◆ Performing the <code>vradmin syncvol</code> command to synchronize volumes to a host running an earlier version of VVR.</li> </ul> <b>Action:</b> Upgrade VVR version on the host to the later version of VVR.
V-5-52-483	Cannot perform the operation: the cross-version feature of VVR does not support the VVR version installed on <i>host</i> .	The cross-version replication in VVR is only supported between two immediate major releases. <b>Action:</b> Upgrade the VVR version on the host <i>host</i> to the same VVR version as on the other hosts in the RDS.
V-5-52-489	Cannot determine local IP address: multiple IP addresses configured for the Primary RVG.	This error occurs when running the <code>vradmin syncrvg</code> command to multiple Secondaries for an RDS in which the <code>local_host</code> field of the Primary RLINKs resolve to different IP addresses. <b>Action:</b> If possible, make the <code>local_host</code> field of all the Primary RLINKs the same. If you still need to configure different interfaces for the Primary RLINKs, then run the <code>vradmin syncrvg</code> command to synchronize one Secondary at a time.



V-5-52-491	Cannot perform the operation: <i>host</i> is not a Primary (acting secondary).	The <code>vradmin fbsync</code> command requires that the specified host be an acting Secondary. <b>Action:</b> Check the <code>vradmin printrvg</code> output to see whether the host <i>host</i> is an acting Secondary.
V-5-52-492	<i>host</i> is an acting secondary. Complete the failback synchronization before making any configuration changes.	When an RDS has an acting Secondary, no other configuration command is allowed on this RDS except the <code>vradmin fbsync</code> command. <b>Action:</b> Run the <code>vradmin fbsync</code> command to convert the acting Secondary to a Secondary before making any configuration changes.
V-5-52-493	Cannot perform the operation: none of the Secondaries is in DCM logging mode.	None of the Primary RLINKs are in DCM logging mode, hence the resync operation cannot be performed. <b>Action:</b> none
V-5-52-494	The command is not available in earlier version of VVR.	This message results when a new command, added in a later VVR version, is run in a cross-version replication environment. <b>Action:</b> To be able to make use of the new commands, make sure that VVR on all hosts in the RDS is upgraded to the appropriate higher version.
V-5-52-502	Host name or IP <i>host</i> is not configured or available on the Primary.	This happens during <code>addsec</code> command, when the host name or IP address specified for the Primary host is not configured or unavailable for use. <b>Action:</b> Make sure the hostname or IP address specified for the Primary host is correct and reachable.



V-5-52-604	Primary RLINK <i>rlk_name</i> not up-to-date.	The command for which this error occurred requires the Primary RLINK <i>rlk_name</i> to be up-to-date. <b>Action:</b> Verify that Primary RLINK <i>rlink_name</i> is up-to-date before running the command. Use <code>vradmin repstatus rvg</code> command or <code>vxrlink status rlk_name</code> command to verify current status of the RLINK.
V-5-52-605	RLINK <i>rlink_name</i> in CONNECT state.	The command for which this error occurred requires that the Primary RLINK <i>rlk_name</i> must not be in CONNECT state. <b>Action:</b> Use <code>vradmin stoprep</code> command to stop replication.
V-5-52-609	Volume <i>vol_name</i> in disk group <i>dg_name</i> not of equal length as Primary volume.	The <code>vradmin addsec</code> and <code>addvol</code> commands require that the Secondary data volume(s) must be of the same size as the corresponding Primary data volume(s). <b>Action:</b> Resize either the Secondary or Primary data volume <i>vol_name</i> to make them of the same size. Then retry running the command. Use the <code>vxresize</code> command to resize a data volume.
V-5-52-610	Primary RLINK <i>rlk_name</i> not in CONNECT state.	The command for which this error occurred requires that the Primary RLINK <i>rlk_name</i> must be in CONNECT state. <b>Action:</b> If the RLINK is not attached, use the <code>vradmin startrep</code> command to start replication. If the RLINK is already attached, but not in CONNECT state, refer to the chapter " <a href="#">Troubleshooting VVR</a> " on page 303 for identifying and correcting the problem.
V-5-52-611	RLINK <i>rlk_name</i> is inconsistent or failed.	This error occurs when the specified Secondary data is inconsistent as a result of which the specific Secondary cannot be converted to Primary. <b>Action:</b> Use another Secondary which is consistent or restore a consistent data from the backup to the Secondary data volumes.



V-5-52-803	Lost connection to host <i>host</i> ; terminating command execution.	<p>During command execution, the host <i>host</i> became unreachable or vradmind server on host <i>host</i> stopped running.</p> <p><b>Action:</b> Retry running the command after host <i>host</i> becomes reachable or after restarting vradmind on host <i>host</i>. This could also be due to some problem with vradmind. In such a case, contact VERITAS Customer Support.</p>
V-5-52-2406	Cannot allocate space to grow volume to <i>new_size</i> blocks.	<p>The system does not have enough free space to accommodate the requested growth in volume size.</p> <p><b>Action:</b> Retry running the command when there is enough free space in VM configuration.</p>
V-5-4-2411	Volumes under RVG <i>rvg_name</i> are not prepared for the instant snapshot.	<p>The command requires the data volumes to be prepared before using the instant snapshot operations.</p> <p>The error occurred when trying to create instant snapshots of all the data volumes under an RVG, because the volumes were not prepared.</p> <p><b>Action:</b></p> <p>For the first time, before creating instant full snapshot of the volumes under an RVG, prepare the volumes using the command:</p> <pre><b>vxsnap -g diskgroup prepare volume</b></pre> <p>For more information, refer to the vxsnap(1M) manual page.</p>

## Messages Related to Configuration Errors

This section describes the configuration related errors that are displayed by the `vradmin -l printrvg` command and the `vradmin -l repstatus` command. The Config Errors section appears in the output of the commands only if there are configuration errors in an RDS. The *host* variable in the error messages is the host name on which the error has occurred.



**Note** Sometimes, the Config Errors section in the `printrvg` command output may display multiple error messages separated by commas, for a specific host.

Message	Message Definition
<p><i>host</i>: Pri or Sec IP not available or vradmind not running</p>	<p>The Primary IP or the Secondary IP address is not available, or the vradmind daemon on the host is not running or is running on a different port.</p> <p><b>Action:</b> If it is a network problem, correct it. If vradmind server is not running on <i>host</i>, start it. If vradmind is running and network connection is fine, make sure that vradmind is using the same port, as vradmind on other hosts in the RDS. Use <code>vrport</code> command to check or set the ports.</p>
<p><i>host</i>: disk group missing.</p>	<p>Host <i>host</i> does not have any disk group with the same name as that specified in the <i>remote_dg</i> attribute, of the Primary RLINK pointing to this host.</p> <p><b>Action:</b> Make sure the <i>remote_dg</i> attribute of the Primary RLINK contains the correct remote disk group name. Use <code>vxprint -l rlink_name</code> to see the RLINK settings. If the disk group needs to be imported on <i>host</i>, and it hasn't been done, then import the disk group on <i>host</i>.</p>
<p><i>host</i>: RLINK missing.</p>	<p>Primary RVG has an RLINK to <i>host</i>, but <i>host</i> does not have corresponding rlink to this Primary RLINK.</p> <p><b>Action:</b> This error can occur due to the following reasons:</p> <ul style="list-style-type: none"> <li>◆ One or more of the following attributes in the Primary RLINK is incorrect: <i>remote_host</i>, <i>remote_dg</i>, and <i>remote_rlink</i>. Make sure these attributes are correct.</li> <li>◆ The corresponding Secondary RLINK on <i>host</i> is missing. To fix this, you can dissociate the Primary RLINK, delete it from the configuration, and then use <code>vradmin addsec</code> command to add <i>host</i> to the RDS.</li> </ul>



Message	Message Definition
<i>host</i> : RLINK dissociated.	Host <i>host</i> does not have an RLINK corresponding to the Primary RLINK. However, it is not associated with the Secondary RVG. <b>Action:</b> Reassociate the RLINK to the Secondary RVG on <i>host</i> .
<i>host</i> : disk-group mismatch.	The <i>remote_dg</i> attribute of either the Primary RLINK or Secondary RLINK is incorrect. <b>Action:</b> Make sure the <i>remote_dg</i> attribute of the Primary RLINK contains the Secondary disk group name and the <i>remote_dg</i> attribute of the Secondary RLINK contains the Primary disk group name.
<i>host</i> : RLINK mismatch.	The <i>remote_rlink</i> attribute of either the Primary RLINK or Secondary RLINK is incorrect. <b>Action:</b> Make sure the <i>remote_rlink</i> attribute of the Primary RLINK contains the Secondary RLINK name and the <i>remote_rlink</i> attribute of the Secondary RLINK contains the Primary RLINK name.
<i>host</i> : host mismatch.	The <i>local_host</i> and/or <i>remote_host</i> attribute of either the Primary RLINK or Secondary RLINK is incorrect. <b>Action:</b> Make sure the value of the <i>local_host</i> attribute of the Primary RLINK is the same as that of the <i>remote_host</i> attribute value of the Secondary RLINK. Also make sure the value of the <i>remote_host</i> attribute of the Primary RLINK is the same as that of the <i>local_host</i> attribute of the Secondary RLINK.
<i>host</i> : Primary-Primary configuration.	The two Primary RVG RLINKs are pointing to each other. This situation will arise after the original Primary comes up after a Primary failover. <b>Action:</b> Use <code>vradmin fbsync</code> or <code>vradmin makesec</code> command to convert the original Primary to a Secondary.



Message	Message Definition
<i>host</i> : multiple Primary error.	<p>The same Secondary RVG has more than one Primary RVGs.</p> <p><b>Action:</b> Check both Primary RVGs to verify which one is the intended Primary then do the following:</p> <ul style="list-style-type: none"> <li>◆ Dissociate the Secondary RLINK pointing to the unwanted Primary RVG.</li> <li>◆ Dissociate the Primary RLINK from the unwanted Primary RVG.</li> <li>◆ Remove the disassociated RLINK(s) from the configuration.</li> </ul>
<i>host</i> : two or more nodes on same host.	<p>Two or more RVGs in the same RDS are located on the same host, <i>host</i>. This configuration is not supported.</p> <p><b>Action:</b> Keep only one RVG on that host and remove the rest.</p>
<i>host</i> : platform mismatch.	<p>Replication is allowed only between hosts that are on the same platform. This is an invalid configuration.</p>
No Primary RVG.	<p>Ignore this error if it displays only for a few seconds and disappears. If the error persists, then the problem may be due to the Secondary not being able to determine its Primary because it has some configuration error, the Primary is not reachable, or it does not have any RLINK.</p> <p><b>Action:</b> If the Secondary has an RLINK to a Primary, check to see if the Primary RVG and the corresponding Primary RLINK exist. If yes, make sure that <code>vradmind</code> is running on Primary and the network between Primary and Secondary is good. If everything is good, then run <code>vradmin -l printrvg</code> command on the Primary to see what type of configuration error is displayed. Correct the configuration error(s).</p> <p>If Primary RVG does not exist or the Secondary RVG does not have any RLINKs, then just remove the Secondary RVG.</p>



Message	Message Definition
<i>host</i> : Primary and Secondary have same disk-group ID.	This condition happens in the cases when the split mirrored plexes of the Primary volumes are exported without using the <b>Disk Group split</b> option and then imported on the Secondary with force option. <b>Action:</b> Contact VERITAS Customer Support.
<i>host</i> : unknown	The configuration status is currently unknown. <b>Action:</b>
<i>host</i> : stale information	The configuration status may be stale. <b>Action:</b>
<i>host</i> : no data volume	This error can occur if one of the RVGs in the RDS does not have a data volume associated with it. <b>Action:</b> Create the appropriate data volume and associate it with the RVG.
<i>host</i> : network-protocol mismatch	The <code>protocol</code> attribute of the Primary RLINK is different from that of the Secondary RLINK. <b>Action:</b> Make sure the value of the <code>protocol</code> attribute for the Primary and Secondary RLINK is the same.
<i>host</i> : VVR-heartbeat-port mismatch	The <code>local_port</code> attribute setting for the Primary RLINK is different from that of the Secondary RLINK. <b>Action:</b> Make sure the value of the <code>local_port</code> attribute for the Primary and Secondary RLINK is the same.
<i>host</i> : unsupported VVR version in cross-version replication	The cross-version replication in VVR is only supported between two immediate major releases. <b>Action:</b> Upgrade the VVR version on the host <i>host</i> to the same VVR version as on the other hosts in the RDS.



Message	Message Definition
<i>host:</i> no contact from Primary	<p>This error can occur if the vradmind server on the Secondary host is unable to establish contact with the vradmind server on the Primary. This can be because the Primary RVG of this RDS cannot be found or the vradmind server is not running or is unreachable.</p> <p><b>Action:</b> If the vradmind server is not running, start it.</p>
<i>host:</i> vxconfigd disabled	<p>The vxconfigd daemon is currently disabled on the host on which this error occurred.</p> <p><b>Action:</b> Start the vxconfigd daemon.</p>
<i>host:</i> volume-number mismatch	<p>This error can occur if the number of volumes in the Primary and Secondary RVGs of the RDS are different.</p> <p><b>Action:</b> Use vradmind printvol to find out which RVG has the extra data volumes and then either remove them from that RVG or associate corresponding data volumes for these extra data volumes to the other RVGs in the RDS.</p>
<i>host:</i> volume-size mismatch	<p>This error can occur if the sizes of some or all of the data volumes in the Primary and Secondary RVGs of the RDS do not match.</p> <p><b>Action:</b> Use vradmind printvol to find out which data volumes have mismatching sizes and then use the vradmind resizevol command to correct the size mismatch error.</p>
<i>host:</i> volume-name mismatch	<p>This error can occur if some or all of the volumes in the Primary and Secondary RVGs of the RDS are not mapped correctly.</p> <p><b>Action:</b> Use vradmind printvol to find out which data volumes are not mapped correctly and correct the mapping error. For more information, refer to the section <a href="#">“Mapping the Name of a Secondary Data Volume to a Differently Named Primary Data Volume”</a> on page 401.</p>



Message	Message Definition
<i>host</i> : Primary SRL missing	This error can occur if the Primary SRL was disassociated from the Primary RVG or is missing. <b>Action:</b> If the Primary SRL is missing, create it and associate it to the Primary RVG. If it was disassociated, re-associate it back to the Primary RVG.
<i>host</i> : Secondary SRL missing	This error can occur if the Secondary SRL was disassociated from the Secondary RVG or is missing. <b>Action:</b> If the Secondary SRL is missing, create it and associate it to the Secondary RVG. If it was disassociated, re-associate it back to the Secondary RVG.

## Messages Related to the vrstat Command

This section lists and describes some of the error messages that may be displayed when running the `vrstat` command. It also suggests the appropriate action that needs to be taken if a problem occurs.

Unique Message Identifier (UMI)	Message	Message Definition
V-5-84-162	Terminating the collection of statistics.	This error may occur while trying to run the <code>vrstat</code> command on the secondary, if the <code>vradmin</code> server on the Primary was terminated while the <code>vrstat</code> command was in progress. This could be due to an administrative action or due to some problem in <code>vradmin</code> . <b>Action:</b> If you had stopped the <code>vradmin</code> server manually, then restart it and run the <code>vrstat</code> command again. If the <code>vradmin</code> server stopped due to some problem, contact VERITAS Customer Support.

# Using the In-Band Control Messaging Utility `vxibc` and the IBC Programming API

---



This appendix explains how to use the IBC Messaging command-line utility `vxibc` and the API for off-host processing. The In-Band Control (IBC) Messaging feature with the Snapshot feature of Volume Replicator (VVR) and optional FastResync (FR) feature of Volume Manager (VxVM) enable you to perform off-host processing. Typically, to perform off-host processing, you would use the `vradm in ibc` command to sequence and automate the operations. For more information, see Chapter 6, “[Using VVR for Off-Host Processing](#)” on page 199. However, if you want to customize the process beyond what can be achieved by using the `vradm in ibc` scripts, or if you want to program and integrate off-host processing in your control facility, you need to use the `vxibc` command or the IBC API.

Off-host processing consists of performing operations on application data on a host other than the one where the application is running. Typical applications include Decision Support Systems (DSS), backup, and trial failover in VVR. In a VVR environment, off-host processing reduces the load on the application server, the Primary, and uses the potentially under-utilized Secondary.

The model for data access on the Secondary is that you break off a mirror from each data volume in the RVG, perform the operation on the mirror, and then reattach the mirror while replication is in progress.



## In-Band Control Messaging Overview

When you take a snapshot on the Secondary, it contains a point-in-time copy of the data on the Primary. Because the Secondary may be behind the Primary, it is not known exactly what time on the Primary this point-in-time copy represents.

VVR maintains a block-level consistency between the Primary and Secondary data volumes. But applications, for example, a file system, that use the data volumes have a higher-level consistency requirement. To support this, VVR provides the IBC facility.

IBC messaging enables you to send a message in the replication stream to notify the Secondary that an event has occurred on the Primary. In the case of a file system, you can use the `sync` command on the Primary, and then send an IBC message. When this message arrives on the Secondary, the data on the Secondary is consistent at the file system level and replication stops. Therefore, further updates are not applied to Secondary data volumes but are stored in the Secondary SRL. You then split off a mirror, which now contains a consistent image of the file system, and unfreeze replication. After the unfreeze operation all the pending updates that are stored in the secondary SRL are applied to the secondary data volumes.

The model with IBC Messaging is that a process on the Secondary waits for the IBC Message and a process on the Primary sends the message when the desired event has occurred.

---

**Note** If you choose not to use IBC Messaging, the data on the Secondary is consistent and can be recovered by the application but it might be out-of-date and potentially out of sync.

---

## Using the IBC Messaging Command-Line Utility

The `vxibc` command-line utility enables you to perform the following IBC Messaging tasks:

- ◆ [Registering an Application Name](#)
- ◆ [Displaying the Registered Application Name for an RVG](#)
- ◆ [Receiving an IBC Message](#)
- ◆ [Sending an IBC Message](#)
- ◆ [Unfreezing the Secondary RVG](#)
- ◆ [Unregistering an Application Name](#)
- ◆ [Receiving and Processing an IBC Message Using a Single Command](#)
- ◆ [Sending and Processing an IBC Message Using a Single Command](#)

For details on using the `vxibc` command and the available options, see the online manual pages.

### Registering an Application Name

Before being able to perform IBC operations on an RVG, you must register an application name for the RVG. The sender and the receivers of the IBC message must register the same application name. Multiple application names (up to a maximum of 32) can be registered for an RVG. Registration is not persistent through host reboots. Applications on rebooted hosts must be reregistered.

To register an application name for an RVG:

```
# vxibc [-g diskgroup] [-D deliver_timeout] \  
      register application_name rvg_name
```

### Displaying the Registered Application Name for an RVG

You can use the `vxibc status` command to display the currently registered application names for a Replicated Volume Group (RVG).

To display the registered application names for an RVG:

```
# vxibc [-g diskgroup] status rvg_name
```



## Receiving an IBC Message

You can use the `vxibc receive` command to receive the IBC message sent from the Primary to a Secondary.

To receive an IBC Message:

```
# vxibc [-g diskgroup] [-n | -R receive_timeout] [-f filename] \  
[-l buf_length] receive application_name rvg_name
```

Note that the `application_name` for the Secondary RVG must have been previously registered.

When the Secondary receives the IBC message, the state of the data volumes on the Secondary is the same as the state of the data volumes on the Primary at the time the IBC message was inserted in the replication stream. Subsequent writes are delivered to the Secondary and stored in the SRL, that is, replication is frozen. Secondary replication remains frozen until an `unfreeze` operation is performed, or the specified `freeze_timeout` expires. The default behavior for the receive operation is to block until an IBC message is received. The option `-n` makes the receive operation non-blocking, and returns if there is no message to receive. If the operation succeeds, the received message is displayed; if a file name is specified the message is written to that file.

## Sending an IBC Message

You can use the `vxibc send` command to send an IBC message from the Primary to a Secondary.

To send an IBC message:

```
# vxibc [-g diskgroup] [-N | -F freeze_timeout] \  
[-f filename | -m message] send application_name rvg_name \  
[rlink_name....]
```

Note that the `application_name` for the Primary RVG must be previously registered.

The IBC message is inserted into the update stream of the specified RLINKs. If an RLINK is not specified, the message is sent to all RLINKs currently attached to the Primary RVG.

IBC messages are always sent to the Secondary RVG irrespective of whether or not the `application_name` is registered on the Secondary.

If the application is registered on the Secondary, then the IBC message is discarded on the Secondary if a receive operation is not performed within the deliver-timeout period.

In the case the application is not registered at the Secondary, then the IBC message is held for ten minutes before being discarded. If the `application_name` is registered within this time, then the IBC message is discarded if a receive operation is not performed within the deliver-timeout period. On the Secondary, the RVG remains frozen until an `unfreeze` operation is performed or the specified `freeze_timeout` expires.

## Unfreezing the Secondary RVG

The `vxibc unfreeze` command unfreezes the Secondary RVG. This operation must be performed after receiving the IBC message using the `receive` operation.

To unfreeze an IBC message:

```
# vxibc [-g diskgroup] unfreeze application_name rvg_name
```

The `vxibc unfreeze` command permits replication to continue by allowing updates that were performed on the Primary data volumes after the `send` operation was executed on the Primary RLINK, to be applied to the Secondary RVG.

## Unregistering an Application Name

The `vxibc unregister` command unregisters an application name for the RVG.

To unregister an application name:

```
# vxibc [-g diskgroup] unregister application_name rvg_name
```

The application name must have been previously registered for the RVG. Further `send` operations against the application name are not possible after unregistering on the Primary RVG.

You can unregister the application on the Secondary if the following conditions are met:

- ◆ If the IBC message has arrived on the Secondary and has been received by the user.
- ◆ If the IBC message has arrived on the Secondary and not received by the user, but the delivery timeout period has expired.

If you used the `vxibc regrecv` command, you do not have to unregister the application.



## Receiving and Processing an IBC Message Using a Single Command

The `vxibc regrecv` command enables you to specify a command to be performed on the arrival of the IBC together with the command arguments. The `vxibc regrecv` command performs the following operations in a single step:

- ◆ Registers the application name
- ◆ Receives the IBC message
- ◆ Runs the specified command with the provided arguments
- ◆ Unfreezes the Secondary RVG
- ◆ Unregisters the application name.

To receive and process an IBC message in a single step:

```
# vxibc [-g diskgroup] [-R receive_timeout] [-f filename] \  
      [-l buf_length] regrecv application_name rvg_name command \  
      [argument]
```





## Sending and Processing an IBC Message Using a Single Command

The `vxibc resend` command performs the following operations in a single step:

- ◆ Registers the application name
- ◆ Sends the IBC message
- ◆ Unregisters the application name

The `vxibc regrecv` command must be started on the Secondary host before the IBC message sent from the Primary host gets invalidated due to a delivery timeout on the Secondary. This can also be done by first executing the `vxibc resend` command on the Primary, followed by the `vxibc regrecv` command on the Secondary within the delivery time-out period which is by default 600 seconds. Otherwise, the IBC message is discarded on the Secondary because there is no corresponding registered application name.

To send and process an IBC message in a single step:

```
# vxibc [-g diskgroup] [-D deliver_timeout] [-N | -F freeze_timeout] \  
[-f filename | -m message] resend application_name rvg_name \  
[rlink_name...]
```

The `vxibc regrecv` command must be issued before the IBC message delivery times out at the Secondary. Typically, this command is issued before the IBC is sent from the Primary.



## Examples—Off-Host Processing

The examples in this chapter assume that the following VVR configuration has been set up on the Primary and Secondary hosts:

Name of the Primary host: `seattle`

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

Name of the Secondary host: `london`

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

The examples use the application name `dss_app` for sending and receiving IBC messages.



For Example 1, Example 2, and Example 3, perform the following steps before you begin:

1. Create a snapshot plex on the Secondary for each data volume using the command:

```
# vxassist -g hrdg snapstart hr_dv01
# vxassist -g hrdg snapstart hr_dv02
```

You can use the `-b` option with the `vxassist snapstart` command to run the command in the background. Note that if you use the `-b` option of the `vxassist snapstart` command, you must wait for the snapshot plexes for all the data volumes in the RVG to be created and synchronized completely. When the plex synchronization completes, the output of the `vxprint` command displays the state of the new snapshot plex as `SNAPDONE`.

2. If you have bought a license for FastResync (FR) use the command:

```
# vxvol -g hrdg set fmr=on hr_dv01
# vxvol -g hrdg set fmr=on hr_dv02
```



## Example 1—Decision Support Using the Traditional Snapshot Feature and the vxibc Utility

1. On the Secondary, register the application name `dss_app` and get ready to receive the IBC message. The command to break-off the snapshot plex when the IBC is received is specified with the `vxibc regrecv` command, as follows:

```
# vxibc -g hrdg regrecv dss_app hr_rvg vxrvrg snapshot hr_rvg
```

2. On the Primary, put the application that is using the Primary data volumes `hr_dv01` and `hr_dv02` into a consistent state.

---

**Note** Putting the application in a consistent state involves flushing all the buffers and pending transactions. For example, the file system can be brought into a consistent state by synchronizing the volumes using the VxFS specific `sync` command. In Oracle, the application can be brought into a consistent state by putting it in a hot-backup mode.

---

3. On the Primary, register the application name `dss_app` and send the IBC message using the following command:

```
# vxibc -g hrdg regsend dss_app hr_rvg
```

When the Secondary receives the IBC message, replication is frozen and no more data is written to the secondary data volumes. The `vxrvrg snapshot` command specified in [step 1](#) then breaks-off the snapshot plexes from the Secondary data volumes, and replication starts again. When the `vxibc` commands complete on both the hosts, the application name is unregistered.

4. On the Primary host, resume the application if it was suspended in [step 2](#).  
If the RLINK is asynchronous and behind there may be a delay between the `vxibc regsend` and `vxrvrg snapshot` command. During this delay, the application is running.
5. On the Secondary, use the snapshot data volumes `SNAP-hr_dv01` and `SNAP-hr_dv02` for running the DSS application, that is, for off-host processing.
6. When the application completes, reattach the snapshot plexes to the data volumes using the following command:

```
# vxrvrg -g hrdg snapback hr_rvg
```

The reattach destroys the SNAP volumes and reattaches the snapshot plexes to their original volumes. If you have enabled FR on these volumes, only the blocks that were changed by the off-host processing application are resynchronized.

## Example 2—Backing Up Using the Snapshot Feature and the vxibc Utility

1. Perform step 1 to step 4 of “[Example 1—Decision Support Using the Traditional Snapshot Feature and the vxibc Utility](#)” on page 380.
2. On the Secondary, copy the snapshot to tapes using a backup utility or the UNIX command `dd`. Use the `dd` command as follows:

```
# dd if=/dev/vx/rdisk/hrdg/SNAP-hr_dv01 of=/dev/rmt/0  
# dd if=/dev/vx/rdisk/hrdg/SNAP-hr_dv02 of=/dev/rmt/0
```

3. Reattach the snapshot plexes to the original volumes:

```
# vxrvg -g hrdg snapback hr_rvg
```



## Example 3—Trial Failover Using the Snapshot Feature

Because the goal is to simulate a crash on the Primary, do not use IBC Messaging for trial failover.

1. Pause the RLINK either on the Primary or the Secondary to maintain consistency. To pause the RLINK on the Primary, type:

```
# vxrlink -g hrdg pause rlk_london_hr_rvg
```

To pause the RLINK on the Secondary, type:

```
# vxrlink -g hrdg pause rlk_seattle_hr_rvg
```

2. When the RLINK is paused, take snapshots of all the data volumes in the RVG:

```
# vxrvrg -g hrdg -P trial snapshot hr_rvg
```

where `trial` is the prefix for the snapshot plexes for all data volumes. Snapshot data volumes with names `trial-hr_dv01` and `trial-hr_dv02` are created.

3. When the snapshots are complete, resume the RLINK by typing:

```
# vxrlink -g hrdg resume rlink_name
```

where `rlink_name` is the name of the paused RLINK.

4. Start the application using the data volumes `trial-hr_dv01` and `trial-hr_dv02` that you have snapped off.
5. Use the recovery function of the application to recover it, and then run the application. For example to recover a `vxfs` file system, use `fsck`.

```
# fsck -F vxfs /dev/vx/rdisk/hrdg/trial-hr_dv01
```

```
# fsck -F vxfs /dev/vx/rdisk/hrdg/trial-hr_dv02
```

6. When the test is complete, shut down the application. For a file system, unmount the file system.
7. Reattach the snapshot plexes to the original data volumes.

```
# vxrvrg -g hrdg -P trial snapback hr_rvg
```

The `-P` option to the `vxrvrg snapback` command reattaches to the original volume the plexes with the prefix specified when taking the snapshot.

## Example 4—Decision Support Using the Instant Full Snapshot Feature and the vxibc Utility

1. On the Secondary, prepare the volumes for which you want to create instant snapshots using the command:

```
# vxsnap -g hrdg prepare hr_dv01
# vxsnap -g hrdg prepare hr_dv02
```

This operation needs to be performed only for the first time you perform the snapshot operation.

2. On the Secondary, create snapshot volumes of the same size as the original volumes and with an appropriate prefix:

```
# vxassist -g hrdg make dss-hr_dv01 volume_length
# vxassist -g hrdg make dss-hr_dv02 volume_length
```

where *volume\_length* is the length of the original volumes.

3. On the Secondary, prepare the snapshot volumes for which you want to create instant snapshots using the command:

```
# vxsnap -g hrdg prepare dss-hr_dv01
# vxsnap -g hrdg prepare dss-hr_dv02
```

4. On the Secondary, issue the following command:

```
# vxibc -g hrdg regrecv dss_app hr_rvg [vxrvrg -g hrdg -F -P dss \
    snapshot hr_rvg]
```

The command `vxrvrg -g hrdg -F -P dss snapshot hr_rvg` is run when the IBC message arrives on the Secondary and the command creates an instant full snapshot.

5. On the Primary, put the application that is using the Primary data volumes `hr_dv01` and `hr_dv02` into consistent state using the application specific method. For information on the consistent state, see [step 2](#) on page 380.

6. On the Primary, register the application name `dss_app` and send the IBC message using the following command:

```
# vxibc -g hrdg regsend dss_app hr_rvg
```

7. On the Primary host, resume the application if it was suspended in [step 5](#).

If the RLINK is asynchronous and behind there may be a delay between the `vxibc regsend` and `vxrvrg snapshot` command. During this delay, the application is running.



8. On the Secondary, use the snapshot data volumes `dss-hr_dv01` and `dss-hr_dv02` for running the DSS application, that is, for off-host processing.
9. When the application completes, reattach the snapshot plexes to the data volumes using the following command:

```
# vxrvg -g hrdg snapback hr_rvg
```

The reattach destroys the `dss` volumes and reattaches the snapshot plexes to their original volumes.

## In-Band Control Messaging API

This section explains how to use the In-Band Control (IBC) Messaging Application Programming Interface (API). VVR supports a special set of ioctls for accessing the IBC messaging facility. These ioctl commands allow an application to register with the facility, send and receive IBC messages, and unregister from the facility.

The IBC facility enables applications to insert application-defined control messages inband with the Primary RVG update stream being replicated to a Secondary RVG. When an IBC message arrives at the Secondary RVG, replication is frozen until directed to unfreeze by a companion application residing on the Secondary host. In this way, an application can signal a Secondary RVG that some user-defined event has occurred relative to the update stream, such as a point of application-level consistency, and enable the Secondary RVG to take some action while replication is frozen.

VVR provides the following ioctl commands:

- ◆ [RV\\_IBC\\_REGISTER](#)
- ◆ [RV\\_IBC\\_SEND](#)
- ◆ [RV\\_IBC\\_RECEIVE](#).
- ◆ [RV\\_IBC\\_UNFREEZE](#)
- ◆ [RV\\_IBC\\_UNREGISTER](#)



## IOCTL Commands

---

**Note** The RVG must be started for the IOCTLs to succeed.

---

RVG devices support five special ioctls: `RV_IBC_REGISTER`, `RV_IBC_UNREGISTER`, `RV_IBC_SEND`, `RV_IBC_RECEIVE`, and `RV_IBC_UNFREEZE`. The format for calling each ioctl command is:

For HP-UX

```
#include <stddef.h>
#include <fcntl.h>
#include <errno.h>
#include <sys/types.h>
typedef uint32_t minor_t;
typedef uint32_t major_t;

#include <vxvm/volmachdep.h>
#include <vxvm/voldefs.h>
#include <vxvm/volioc1.h>
#include <vxvm/volibc.h>
int ioctl(int fd, int cmd, void *arg);
```

The argument `fd` is the file descriptor obtained by opening the RVG device using the `open(2)` system call.

The value of `cmd` is the ioctl command code, and `arg` is a pointer to a structure containing the arguments to be passed to the kernel. Definitions of the argument structures for each ioctl are described below.

The return value for all ioctls is 0 if the command was successful, and -1 if it was rejected. If the return value is -1, then `errno` is set to indicate the cause of the error.

### RV\_IBC\_REGISTER

This ioctl registers an application name for the RVG and returns a key. Only registered application names, using the key, may use the IBC messaging facility on a particular RVG. Multiple application names can be registered for any RVG, up to a maximum of 32.

The ioctl argument structure for the `RV_IBC_REGISTER` command is:

```
struct ibc_register_args {
    char          application_name[NAME_SZ];
    int           deliver_timeout;
    ibc_appid_t   application_id;
};
```



Argument `deliver_timeout` specifies a time-out value in seconds for delivery of an IBC message after it has arrived at the Secondary RVG. When the time-out expires, the Secondary RVG discards the IBC message and continues replication. See `RV_IBC_SEND` and `RV_IBC_RECEIVE` for definition of message delivery. A `deliver_timeout` of 0 is used to specify no time-out.

Argument `application_id` is returned by the `ioctl`. It must be supplied as input argument to all other IBC `ioctls`.

Use of IBC messages is inherently distributed. A copy or agent of the application is expected to be resident on each participating host, and each participating application must register on its own host. Those resident on the Secondary host must register using an application name identical to the name registered on the Primary host. The returned *application\_id* has a local scope; it can be distributed to any cooperating applications on the same host, but it cannot be used successfully by an application on a remote host.

An IBC message received on the Secondary for an application name that is not registered is discarded after delivery timeout. Registration is not persistent across system reboots. Applications must be registered again after the host reboots. After the Secondary is rebooted, the application must be registered within ten minutes after `vxnetd` is started if an IBC message has already arrived.

The `vxnetd` command is started from the system startup script:

```
/sbin/rc2.d/S994vxnm-vxnetd
```

On failure, possible values returned in `errno` are:

<code>EIBC_NOMEM</code>	Maximum number of applications (32) already registered.
<code>EIBC_DUP_APPLICATION</code>	<i>application_name</i> is already registered.

## RV\_IBC\_SEND

This ioctl can only be issued against the Primary RVG with a valid key obtained from the `RV_IBC_REGISTER` ioctl. The ioctl inserts an IBC message into the data update stream of one or all RLINKs attached to the RVG.

If it is desired that an IBC message be inserted at an exact location in the update stream, such as a point of application-level consistency, then there must be no concurrent write activity to the RVG when the `RV_IBC_SEND` ioctl is issued. Note that writes made to the block device interface of a data volume may be cached, so a disk sync must be done before issuing the ioctl. If there are active writes to the RVG when the ioctl is issued, the insertion point of the IBC message in the RLINK update data stream is arbitrary in relation to that activity.

The ioctl returns using the same semantics as a data write to the RVG; it returns when the IBC message has been committed to the SRL and has also been transferred to all synchronous-mode replicas attached to the RVG.

The ioctl argument structure for the `RV_IBC_SEND` command is:

```

struct ibc_send_args {
    vx_u32_t    ibc_magic;
    vx_u32_t    ibc_version;
    ibc_appid_t application_id;
    char        replica[NAME_SZ];
    int         flags;
    int         freeze_timeout;
    caddr_t     msg_buf;
    int         msg_len;
};

```

Argument *ibc\_magic* is used to verify whether the ioctl structure is a valid 4.0 structure. It should be set to `NM_IBC_MAGIC`.

Argument *ibc\_version* specifies the current IBC version. It should be set to `NM_IBC_VERSION`.

Argument *application\_id* is the key returned by the `RV_IBC_REGISTER` ioctl. A registration must be done before the `RV_IBC_SEND` ioctl can be used.

Argument *replica* specifies the name of the RLINK to which the IBC message is to be send. The null string specifies a broadcast to all RLINKs currently attached to the Primary RVG.

Argument *flags* set to `RV_IBC_FREEZE` causes the secondary replication to freeze for the time-out period specified in *freeze\_timeout*. If replication is not desired to be frozen, then *flags* should be set to 0.



Argument *freeze\_timeout* specifies a time-out value in seconds between delivery of an IBC message on the Secondary and execution of an `RV_IBC_UNFREEZE` ioctl against the Secondary RVG. When the time-out expires, replication at the Secondary continues. A time-out value of zero is used to specify no time-out.

Argument *msg\_buf* is a pointer to a buffer containing an IBC message. The content of an IBC message is user-defined and has no restriction except size.

Argument *msg\_len* is the length, in bytes, of the IBC message and can be no greater than 128k bytes.

On failure, possible values returned in `errno` are:

<code>EIBC_NO_RLINK</code>	No <i>RLINK</i> or specified <i>RLINK</i> exists
<code>EIO I/O</code>	I/O error while logging the IBC message
<code>EIBC_MESSAGE_LENGTH</code>	Message is greater than maximum allowable length (128K)

## RV\_IBC\_RECEIVE

This ioctl can only be issued against a Secondary RVG with a valid key obtained from the `RV_IBC_REGISTER` ioctl. The ioctl receives an IBC message sent from the Primary RVG. At the time of receipt, Secondary replication is frozen. The state of the data volumes on the Secondary is the same as that on the Primary at the time the IBC message was sent. Secondary replication remains frozen until an `RV_IBC_UNFREEZE` ioctl is issued against the Secondary RVG, or the *freeze\_timeout* specified when the IBC message was sent expires, or the *deliver\_timeout* specified when the application name was registered for the Primary RVG expires and the receive operation has not been performed.

The ioctl argument structure for the `RV_IBC_RECEIVE` command is:

```
struct ibc_receive_args {
    ibc_appid_t    application_id;
    int           flags;
    ibc_timeout_t  timeout;
    int           drop_count;
    caddr_t       msg_buf;
    size_t        buf_len;
    size_t        msg_len;
};
```

Argument *application\_id* is the key returned by the `RV_IBC_REGISTER` ioctl. A registration must be done before the `RV_IBC_RECEIVE` ioctl can be used.

Argument *flags* may specify `IBC_BLOCK`. If this flag is set, the `ioctl` will block until an IBC message is available to receive. If `IBC_BLOCK` is not set, the `ioctl` returns with an error if no IBC message is available.

Argument *timeout* specifies a time-out value in seconds to block waiting for an IBC message if flag `IBC_BLOCK` is set. When the time-out has expired, the `ioctl` returns with an error. A time-out of zero is used to specify no time-out.

Value *drop\_count* is returned by the `ioctl`. This value contains the number of messages that have been dropped due to delivery time-outs. If *drop\_count* is non-zero, no message is returned and the `ioctl` returns an error.

Argument *msg\_buf* is a pointer to a buffer to receive the IBC message.

Argument *buf\_len* is the length, in bytes, of the *msg\_buf*.

Value *msg\_len* is returned by the `ioctl` and specifies the length in bytes of the IBC message. The maximum IBC message length is 128K bytes. If *msg\_len* is greater than *buf\_len*, the IBC message is truncated to *buf\_len* bytes and no error is indicated.

On failure, possible values returned in `errno` are:

`EIBC_NO_APPLICATION` Argument *application\_id* is not valid.

`ENOMSG` IBC messages have been dropped due to delivery time-out, or if no IBC message was available.



## RV\_IBC\_UNFREEZE

This ioctl can only be issued against a Secondary RVG with a valid key obtained from the `RV_IBC_REGISTER` ioctl. The ioctl unfreezes replication of the Secondary RVG; that is, it resumes updates of the Secondary volumes.

The ioctl argument structure for the `RV_IBC_UNFREEZE` command is:

```
struct ibc_unfreeze_args {
    ibc_appid_t    application_id;
};
```

Argument `application_id` is the key returned by the `RV_IBC_REGISTER` ioctl. A registration must be done before the `RV_IBC_UNFREEZE` ioctl can be used.

On failure, the possible values returned in `errno` are:

`EIBC_NO_APPLICATION` Argument *application\_id* is not valid.

`EBUSY` There is currently active use of the IBC messaging facility by one or more ioctls using this *application\_id*

## RV\_IBC\_UNREGISTER

This ioctl unregisters an application name. This ioctl returns an error if any ioctl is active for the RVG using the registered key.

For a Primary RVG, no `RV_IBC_SEND` ioctls will be accepted with the registered key after unregistering. Those IBC messages already introduced into the update stream are not affected by a subsequent unregister, even if they have not yet been sent to the Secondary RVG.

For a Secondary RVG, `RV_IBC_RECEIVE` or `RV_IBC_UNFREEZE` ioctls using the registered key cannot be successfully executed after the unregister, and any IBC messages arriving for the registered name are discarded.

The ioctl argument structure for the `RV_IBC_UNREGISTER` command is:

```
struct ibc_unregister_args {
    ibc_appid_t    application_id;
};
```

Argument *application\_id* is the key returned by the `RV_IBC_REGISTER` ioctl. A registration must be done before the `RV_IBC_UNREGISTER` ioctl can be used.

On failure, possible values returned in `errno` are:

`EIBC_NO_APPLICATION` Application is not registered.

`EBUSY` IBC deliver or unfreeze pending.

## Using the IBC API

The `ioctl` command set is intended to be used by a set of daemons, one on the RVG Primary host and one on each Secondary host that is to participate in IBC message retrieval. Each must register under an identical application name and be registered before IBC message generation begins. Because registration does not survive host crashes, but IBC messages once sent do persist beyond host crashes, it is suggested that the Secondary daemons be spawned as a part of system startup.

IBC messages use at-least-once delivery semantics. Retrieval daemons must be tolerant of receiving the same IBC message more than once. It is however guaranteed any duplicate copies of a messages will be delivered before the next new message is delivered.







# VERITAS Volume Replicator Object States



This appendix describes the Kernel State (KSTATE) and utility state (STATE) of the RVG and RLINK objects.

---

**Tip** In most cases, if the KSTATE/STATE is *ENABLED/ACTIVE*, the object is available.

---

## VERITAS Volume Replicator Kernel State

The Kernel State (KSTATE) indicates the accessibility of the RVG and RLINK objects. When you issue the `vxprint` command, the KSTATE is listed under the KSTATE heading of the output.

### RVG KSTATES

- ◆ **ENABLED**—You can do I/O to the volumes that belong to the RVG.
- ◆ **DISABLED**—You cannot do I/O to the volumes that belong to the RVG. It must be enabled before you can use the RVG. Issue the `vxrvg start` command.
- ◆ **RECOVER**— You cannot do I/O to the volumes that belong to the RVG. This state is triggered after a diskgroup import or if the RVG object was not recovered properly after a reboot or a crash. Issue the `vxrvg recover` command.

### RLINK KSTATES

- ◆ **CONNECT**—Replication is taking place.
- ◆ **ENABLED**—The RLINK is not communicating with its peer; therefore, replication is not taking place. When the RLINK does communicate with its peer, it automatically reverts to the **CONNECT** state.
- ◆ **DETACHED**—The RLINK is not replicating and is not attempting to connect. Issue the `vxrlink att` command.



- ◆ RECOVER—The RLINK is out of operation. This state is triggered after a diskgroup import or if the RLINK object was not recovered properly after a reboot or a crash. Issue the `vxrlink recover` command.

# VERITAS Volume Replicator Utility States

## RVG Utility States

This section lists the RVG states and their descriptions.

- ◆ **EMPTY**—State of a newly created RVG. Issue the `vxrvrg start` command to start the RVG.
- ◆ **CLEAN**—The RVG is stopped. This state is seen after issuing the `vxrvrg stop` command. Issue the `vxrvrg start` command to start the RVG.
- ◆ **ACTIVE**—This determines if I/O can be performed to the data volumes:
  - ◆ If the **KSTATE** is **ENABLED**, I/O can be performed.
  - ◆ If the **KSTATE** is **RECOVER**, I/O cannot be performed (this state usually occurs after a system crash).
  - ◆ If the **KSTATE** is **DISABLED**, I/O cannot be performed.
- ◆ **FAIL**—A data volume error occurred.

## RLINK Utility States

This section lists the RLINK states and their descriptions.

- ◆ **UNASSOC**—Not associated to an RVG.
- ◆ **STALE**—Associated to an RVG but needs complete synchronization between the Primary and Secondary.
- ◆ **ACTIVE**—Replicating or ready to replicate.
- ◆ **PAUSE**—Replication is not active because of an administrative action or a configuration error.
- ◆ **FAIL**—Data volume error occurred on the Secondary or the `vxrlink -w pause` command was issued on the Secondary. For more information, see [“Inconsistent RLINKs”](#) on page 398.
- ◆ **PAUSING**—Temporary state while `vxrlink pause` is executing.
- ◆ **RESUMING**—Temporary state while `vxrlink resume` is executing.
- ◆ **RESTORING**—Temporary state while `vxrlink restore` is executing.



## Inactive RLINKs

An RLINK is considered inactive whenever the Primary cannot send data to the Secondary for any reason, such as:

- ◆ A temporary failure of the network connection between the Primary and the Secondary.
- ◆ A failure of the Secondary node.
- ◆ The execution of a `vxrlink pause` command by an administrator.

The data to be sent on an inactive RLINK is buffered in the SRL. If the RLINK remains inactive for an extended period, the SRL may not be able to buffer new writes; even if it can, the Secondary becomes increasingly out-of-date. So it is important that the SRL is large enough to accommodate most of the inactive periods. To help control the behavior of an inactive RLINK, SRL overflow protection may be used. For details, see [“Setting the SRL Overflow Protection”](#) on page 53.

## STALE RLINK State

An RLINK is STALE when the Secondary data volumes do not contain the Primary’s data and cannot be brought up-to-date using the SRL. When an RLINK is first created, its initial state is STALE.

RLINKs can enter the STALE state when they are detached either manually (via `vxrlink det`) or by the kernel (on the Primary only, if an SRL media error occurs). An RLINK can also become STALE if the log overflows. See [“Protection Against SRL Overflow—`srlprot` attribute”](#) on page 28 for a description of cases where the log overflows and what can be done to prevent it.

To change the state of an RLINK from STALE to ACTIVE, see [“Using the Automatic Synchronization Feature”](#) on page 61 and [“Example—Synchronizing the Secondary Using Block-level Backup”](#) on page 67.

## FAIL RLINK State

A Primary RLINK enters the FAIL state when the corresponding Secondary RLINK enters the FAIL state for any reason. This happens if there is an unrecoverable I/O error on one of the Secondary data volumes.

There are two ways for a Secondary RLINK to fail. One is if it encounters an I/O error that cannot be corrected. This is less likely to happen if the data volumes have been configured in a redundant fashion.

The second way for a Secondary RLINK to fail is if you enter the `vxrlink -w pause` command on the Secondary. This command must be used with great care because it enables data volumes on the Secondary to be written. The command must be used if a data volume must be restored from backup.

When the restore operation is complete, execute the following command:

```
# vxrlink -g diskgroup -c checkpoint_name restore rlink_name
```

This will return both the Primary and Secondary RLINKs to the ACTIVE state.

Secondaries can also be restored from a Primary checkpoint if a Secondary checkpoint is not available, but the Primary checkpoint and corresponding backup are available.

If the Secondary RLINK cannot be restored, or if it is no longer needed, then `vxrlink det` can be used on either the Primary or the Secondary to detach the Primary RLINK and make it STALE.

---

**Note** In some cases after an RLINK has been moved from the FAIL state back to the ACTIVE state, the RVG may remain in the FAIL state. This can be corrected by entering:

```
# vxrvg -g diskgroup start rvg_name
```

---



## Inconsistent RLINKs

When an RLINK is inconsistent, the Secondary cannot be used for a failover because the data volumes do not reflect the data on the Primary node.

Note that `inconsistent` is not a state. Whether an RLINK is consistent, or not, is shown in the `flags` field of the RLINK.

To see whether the `consistent` or `inconsistent` flag is set, use the following command:

```
# vxprint -g diskgroup -l rlink_name
```

An RLINK that is `inconsistent` and in the `FAIL` state must be restored before it can be used again. It can be restored using a Primary or a Secondary checkpoint. An RLINK becomes `inconsistent` and gets in the `FAIL` state in situations such as:

- ◆ When you enter the `vxrlink -w pause` command  
Used to put an RLINK in the fail state. Not normally used.
- ◆ If there is an unrecoverable I/O error on a data volume on the Secondary  
If the data volume can be restored from backup, it is possible to recover. Loss of volumes due to I/O errors is usually preventable by mirroring.

If an RLINK is `inconsistent`, but not in the `FAIL` state, it could be a temporary situation and the `inconsistent` flag will clear when the operation completes. This happens in situations, such as:

- ◆ During atomic update  
An atomic update operation would happen automatically, for example, to catch up after a network outage. If a machine crashed during such an update, the user would see the `inconsistent` flag set while not in the `FAIL` state. This is unlikely, however, and as long as the Primary node has not been lost, VVR will automatically make the RLINK consistent again once the Primary-Secondary network connection is reestablished.
- ◆ During DCM resynchronization  
When you execute the `vxrvg resync` command after the SRL has overflowed, the RLINK becomes `inconsistent` during resynchronization until the DCM replay is complete.

When the `inconsistent` flag is set, a flag is displayed indicating whether the RLINK can be resynchronized. If the RLINK has the `cant_sync` flag set, it is inconsistent, and this Secondary needs to be resynchronized before it can take part in replication again. If the `inconsistent` and `can_sync` flags are set, there is enough information to make it consistent again. This will occur automatically.

## Pausing, Resuming, and Restoring RLINK States

PAUSING, RESUMING, and RESTORING are temporary states through which the RLINK transitions when doing a `Pause`, `Resume`, or `Restore`, respectively. If these states persist, it means that the command failed halfway through the execution. Recovery from these states is simple.

If the state is PAUSING, it means that some error prevented the pause operation from completing. The error is displayed during execution of the `vxrlink pause` command. When the error is corrected, the next `vxrlink pause` command will succeed.

If the state is RESUMING, it means that some error prevented the resume operation from completing. The error is displayed during execution of the `vxrlink resume` command. When the error is corrected, the next `vxrlink resume` command will succeed.

If the state is RESTORING, a `vxrlink restore` command failed. You must execute either a `vxrlink -w pause` command to put the RLINK back into the FAIL state, or a `vxrlink -c checkpoint restore` command to put it into the ACTIVE state.

Two other Volume Replicator commands also use two-phase transactions. If these commands fail after executing partially, they can be safely repeated. The commands are:

- ◆ `vxrlink recover`
- ◆ `vrxvrg recover`

If the `vxrlink recover` or `vrxvrg recover` command fails, the state of the object will still be RECOVER. Volume Replicator commands that do two-phase transactions report an error message and have a nonzero exit code if they fail.







## VVR Task Reference

---

This appendix describes how to perform the following tasks:

- ✓ [Mapping the Name of a Secondary Data Volume to a Differently Named Primary Data Volume](#)
- ✓ [Mapping Disk Groups](#)
- ✓ [Decreasing the Size of the SRL on the Primary](#)

### Mapping the Name of a Secondary Data Volume to a Differently Named Primary Data Volume

We recommend that you use the same name for a data volume in a Primary RVG and the corresponding data volume in a Secondary RVG. However, each Secondary data volume can have a different name from that of the corresponding Primary data volume. The Primary does not know whether the name is mapped to a different name at any given Secondary. The name mapping information is maintained entirely at the Secondary. To facilitate name-mapping, each data volume associated to an RVG has a `primary_datavol` field. This field can be set to the name of the corresponding data volume on the Primary.

By default, *global mapping* is in effect, that is, the `primary_datavol` field is not used on any Secondary data volumes. This requires that all the Secondary data volumes have the same names as used on the Primary.

One of the prerequisites for adding a Secondary using the `vradmin addsec` command is that data volumes of the same names and lengths as the Primary must exist on the Secondary. When adding a Secondary using the `vradmin addsec` command, the Secondary data volumes cannot have different names from that of the corresponding Primary data volumes.



If you use different names for the Secondary data volumes and their corresponding Primary data volumes, the `vradm migrate` command does not set the `primary_datavol` field on the new Primary after transferring the Primary role. To facilitate seamless transfer of the Primary role, make sure you set the `primary_datavol` field of the Primary data volumes, in addition to the Secondary. Note that you can use the `vradm` command to perform all other VVR operations in a configuration containing differently named volumes on the Primary and Secondary.

There are two ways to set the `primary_datavol` field on a Secondary data volume. In the examples that follow, the commands are executed only on the Secondary. The Secondary data volume is called `secondaryname-dv_name`, and the corresponding Primary data volume name is `dv_name`.

1. To set the name of the corresponding Primary data volume on a Secondary data volume after the volume has been associated to the Secondary RVG, use the `vxedit` command:

```
# vxedit -g diskgroup set primary_datavol=dv_name \  
secondaryname-dv_name
```

OR

1. To set the name of the Primary data volume on a corresponding Secondary data volume when it is being associated with the Secondary RVG, specify the `-m` option on the `vxvol` command line:

```
# vxvol -g diskgroup -m assoc rvg_name \  
secondaryname-dv_name dv_name
```

2. On the Secondary, display the `primary_datavol` field for a volume using `vxprint -l`:

```
# vxprint -g diskgroup -l secondaryname-dv_name
```

### Output resembles:

```
Volume: secondaryname-vol03  
assoc: rvg=rvg_name  
plexes=secondaryname-vol03-01  
primary_datavol=dv_name
```

---

**Note** If any volume (on the Primary or a Secondary) is associated with an RVG (as an SRL or a data volume), the `vxprint -l` listing for that volume will indicate the RVG name on the output line beginning with `assoc`; as shown above.

---

## Mapping Disk Groups

If the RVGs on the Primary and Secondary are in differently named disk groups, the disk group mapping can be specified either when the RLINK is created, or later.

For example, if the disk group on the Primary is *dg1* and on the Secondary is *dg2*, then use the following commands to map the disk group during RLINK creation:

Primary:

```
# vxmake -g dg1 rlink rlink_name remote_dg=dg2
```

Secondary:

```
# vxmake -g dg2 rlink rlink_name remote_dg=dg1
```

If the disk groups were not properly mapped at the time of RLINK creation, the RLINK cannot be attached. This problem can be corrected as follows:

Primary:

```
# vxedit -g dg1 set remote_dg=dg2 rlink_name
```

Secondary:

```
# vxedit -g dg2 set remote_dg=dg1 rlink_name
```



## Decreasing the Size of the SRL on the Primary

Before resizing the SRL, do the following:

- ✓ Stop the application.
- ✓ Verify that the RLINKs are up-to-date by issuing the following command:

```
# vxrlink -g diskgroup status rlink_name
```

### ▼ To decrease the size of the SRL on the Primary

1. If the application and the RVG are *not* configured as VCS resources, proceed to the next step.

OR

If the application and the RVG are configured as VCS resources, OFFLINE the application resource, as well as, the RVG resource and then proceed to [step 4](#). To OFFLINE resources, use the `hagrps` command. For more information on offlining resources, see the *VERITAS Cluster Server User's Guide*.

2. Make sure that the application is not running.

3. Stop the RVG:

```
# vxrvrg -g diskgroup stop rvg_name
```

4. Make sure all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

---

**Note** If you see any outstanding writes, do not proceed to [step 5](#).

---

5. Detach all RLINKs:

```
# vxrlink -g diskgroup det rlink_name
```

6. Dissociate the SRL from the RVG.

---

**Note** Any checkpoints that you have created will be lost after dissociating the SRL.

---

```
# vxvol -g diskgroup dis srl_name
```

7. Decrease the size of the Primary SRL using the `vxassist` command. For example, to decrease the size of the SRL:

```
# vxassist -g diskgroup shrinkto srl_name new_length
```

**Note** It is recommended that the SRL reside on disks that are not being used for the data volumes. Also, it is recommended that the Primary and Secondary SRLs must be of the same size. See the `vxassist (1M)` manual page for more information.

---

8. Reassociate the SRL with the RVG:

```
# vxvol -g diskgroup aslog rvg_name srl_name
```

9. Attach all RLINKs:

```
# vxrlink -f att rlink_name
```

---

**Note** The RLINK was up-to-date when it was detached in [step 5](#) on page 404 and the Secondary is consistent with the Primary; therefore, it is appropriate to use the force option here.

---

10. If the application and the RVG are *not* configured as VCS resources, proceed to the next step.

OR

If the application and the RVG are configured as VCS resources, `ONLINE` the RVG resource, as well as, the application resource. To `ONLINE` resources, use the `hagrps` command. For more information on onlining resources, see the *VERITAS Cluster Server User's Guide*. The resize operation is complete. Do not perform [step 11](#) and [step 12](#).

11. Start the RVG:

```
# vxrvg -g diskgroup start rvg_name
```

12. Restart the application.





## Accessibility and VVR VEA

---

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 also are provided as accessible PDF files, and the online help is provided as HTML displayed in a compliant viewer.

The following topics detail accessibility features and compliance in VERITAS Enterprise Administrator for VERITAS Volume Replicator (VVR VEA):

- ◆ [“Keyboard Navigation and Shortcuts in VVR VEA”](#) on page 408
- ◆ [“Support for Accessibility Settings and Assistive Technologies”](#) on page 415



## Keyboard Navigation and Shortcuts in VVR VEA

All program functions and menu items are accessible using the keyboard exclusively. VVR VEA uses standard operating system navigation keys and keyboard shortcuts. For its unique functions, VVR VEA uses its own navigation keys and keyboard shortcuts which are documented below.

To see a table of the standard Microsoft navigation keys and keyboard shortcuts, select your version of Microsoft Windows from the drop-down listbox at:

- ◆ <http://www.microsoft.com/enable/products/keyboard/keyboardsearch.asp>

Exceptions and additions to the Microsoft Keyboard Conventions are listed in the following tables:

- ◆ [VEA Exceptions to Microsoft Keyboard Conventions](#)
- ◆ [VEA Additions to Microsoft Keyboard Conventions](#)
- ◆ [VEA Help Additions to Microsoft Keyboard Conventions](#)

### VEA Exceptions to Microsoft Keyboard Conventions

VEA Exceptions to Microsoft Keyboard Conventions

Convention	Mnemonic	Accelerator
Cancel	C	Alt C, Escape
OK	O	Alt O
Refresh	E	None
Yes	None	None
No	None	Escape
Shift+F10	A	Alt A



## VEA Additions to Microsoft Keyboard Conventions

### VEA Additions to Microsoft Keyboard Conventions

Convention	Mnemonic	Accelerator
<u>A</u> ctions	A	Alt A. Use this accelerator in place of Shift+F10 with a selected item to bring up the context menu with the Actions menu.
<u>A</u> dd category	A	None
<u>A</u> dd	D	Alt D
<u>C</u> onnect	C	None
<u>C</u> ontents	C	None
Delete Now	None	None
<u>D</u> elete Temporary files on Exit	E	Alt E
<u>D</u> isconnect	D	None
<u>N</u> umber of items in history	N	Alt N
<u>P</u> assword	P	Alt P
<u>P</u> revious	P	Alt P
Over <u>w</u> rite the log file	W	Alt W
Overwrite the log file (s) <u>a</u> fter	A	Alt A
<u>O</u> verwrite the log file (s) older than	O	Alt O
<u>R</u> emove	E	Alt E
Re <u>s</u> can	S	None
<u>R</u> eset	R	Alt R
<u>S</u> et to Defaults	S	Alt S
<u>R</u> emember Password	R	Alt R
<u>U</u> sername	U	Alt U



<u>M</u> anage Profiles	M	Alt T, M
<u>W</u> indow	W	Alt W
<u>T</u> ear off	T	Alt W, T
<u>E</u> nter a name for this profile	E	Alt E
<u>P</u> assword protect this profile	P	Alt P
Enter <u>P</u> assword	A	Alt A
<u>V</u> erify new password	V	Alt V
<u>P</u> assword	A	Alt A
<u>S</u> elect Profile	S	Alt S
<u>P</u> assword	P	Alt P
<u>M</u> anage Profiles	M	Alt M
<u>P</u> rofile Names	R	Alt R
<u>S</u> tart up using	S	Alt S
<u>N</u> ew	N	Alt N
<u>E</u> dit	E	Alt E
<u>D</u> elete	D	Alt D
<u>C</u> lose	O	Alt O
<u>O</u> ld Password	L	Alt L
<u>N</u> ew Password	N	Alt N
<u>V</u> erify new password	V	Alt V

## VEA Help Additions to Microsoft Keyboard Conventions

VEA Applications/Help Additions to Microsoft Keyboard Conventions

Convention	Mnemonic	Accelerator
<u>F</u> ind in Topic...	F	Ctrl + F
Find <u>N</u> ext	N	F3
Find <u>P</u> revious	P	None
<u>H</u> ide Navigation Tabs	H	None
<u>I</u> ndex	I	None
<u>P</u> rint	P	None
<u>S</u> earch	S	None
Synchron <u>u</u> ze	N	None

## General Keyboard Navigation Within the GUI

You can navigate and use VVR VEA with only the keyboard. In the GUI, the current active tab, radio button, or checkbox is enclosed within a rectangle formed by dotted lines. These areas are said to have *focus* and will respond to commands.

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 screen.



## Keyboard Navigation Within Dialog Boxes

Dialog boxes contain groups of controls necessary to set options or settings for programs. Here are some general rules about dialog box navigation:

- ◆ Tab moves focus between controls within the dialog box along a preset sequence.
- ◆ Controls displaying a mnemonic (an underlined letter) can be selected regardless of focus by typing Alt and the underlined letter.
- ◆ A dark border indicates the default command button. Press Enter at any time to choose the button with a dark border.
- ◆ Spacebar chooses a control you select with Tab.
- ◆ Spacebar changes the state of a checkbox or radio button that has focus. Typing a mnemonic (if one is available) will move the focus to the checkbox or radio button and change its state.
- ◆ Arrow keys move focus within listboxes, groups of option controls, or groups of page tabs.
- ◆ Items that cannot be changed are not visited by the Tab key sequence. Options that are unavailable are grayed-out and can neither be selected nor given focus.

While the controls described here are typically found in dialog boxes, they also can occur in other contexts. The same navigation standards will apply.

**Command buttons** The **OK** button carries out the command using the information supplied in the dialog box. Other command buttons enables you to cancel the command or choose from additional options.

Command buttons marked with an ellipsis (...) open another dialog box so you can provide more information or confirm an action.

A dark border initially indicates the default button. Press ENTER at any time to choose the button with a dark border. Use TAB to move the keyboard focus to the next control, and add SHIFT to reverse the direction. When you move the keyboard focus to a command button it temporarily takes the dark border, but when the keyboard focus is not on a command button the dark border returns to the default button. Press SPACEBAR to choose the command button you selected with TAB. Use ESC to immediately choose the Cancel button if there is one. Typing the access key for a command button will immediately carry out the command.

**Check boxes** Use TAB to move between boxes and SPACE to change the check box to the next state. Typing the access key for a check box also will move the focus to the box and change its state.

**Option controls** (also known as radio buttons) are used to select only one option from a group of options. (Like check boxes, option buttons may represent two or three states.) In some applications, you can add CTRL to move to the button without selecting it, then use SPACEBAR to choose the currently selected option button. Typing the access key for an option control will also move the focus to the control and select it.

**Tabbed pages** are used to fit many options into a single dialog box. They appear visually similar to dividers from a notebook. Each page contains separate groups of controls such as check boxes or option controls. Use TAB to move the focus to the tab for the currently visible page. Typing the access key for a page tab will also move the focus to the tab and display its page.

#### Tabbed Pages

Keys:	Action:
CTRL+PAGE DOWN OR CTRL+TAB	Switches to the next tab and displays the page.
CTRL+PAGE UP	Switches to the previous tab and displays the page.
Right arrow or left arrow	When the focus is on a tab selector, chooses the next or previous tab in the current row and displays the page.
DOWN ARROW or UP ARROW	Chooses the tab in the next or previous row and displays the page.

**List boxes** display a column of available choices. If there are more choices than can fit in the open list box, moving the selection through the list will scroll more choices into view. Use TAB or type the appropriate access key to move the focus to the list box.

**Single selection list boxes** allow only one item to be selected from the list at a time. Use the UP ARROW or DOWN ARROW keys to select items. Moving to a second item will automatically deselect all other items.

**Drop-down list boxes** by default show only the selected item. There is a small button to the right of the control that shows a downward-pointing arrow. You can click the arrow to display more items from the list. You also can show or hide the list using ALT+DOWN ARROW, ALT+UP ARROW, or F4. In all other respects it behaves like a normal single-selection list box.



Combo-boxes and drop-down combo boxes are list boxes with attached edit boxes. You can make your selection by either typing in the edit box or using the standard commands to select from the list.

### Combo Boxes and Drop-Down Combo Boxes

Keys	Action
UP ARROW or DOWN ARROW	Selects the previous or next item.
PAGE UP or PAGE DOWN	Selects the item up or down one screen.

**Edit controls** are rectangular boxes used for displaying or editing text. Use TAB or the appropriate access key to move the focus to the edit control and select the text.

## Support for Accessibility Settings and Assistive Technologies

VERITAS software responds to operating system accessibility settings.

VERITAS products are compatible with Microsoft's accessibility utilities. In Windows 2000, accessibility options involving keyboard responsiveness, display contrast, alert sounds, and mouse operation can be set through the Control Panel (**Start > Settings > Control Panel > Accessibility Options**) and through the Accessibility Wizard (**Start > Programs > Accessories > Accessibility > Accessibility Wizard**).

Also, VERITAS software is compatible with many third-party assistive technologies. VVR VEA is compatible with JAWS.

---

**Note** Though all graphics in VERITAS documentation can be read by screen readers, setting your screen reader to ignore graphics may improve performance.

---







# Glossary

---

## **ACTIVE**

On a Secondary RLINK, the **ACTIVE** state indicates that it is ready to receive updates from the Primary.

## **asynchronous**

Asynchronous mode queues writes persistently and holds them at the Primary for later transmission. This mode can deal with temporary outages of the network or the Secondary node without impacting the application. Asynchronous mode is useful when it is acceptable for the Secondary not to be up-to-date.

## **Automatic Synchronization**

A feature of VVR that automatically synchronizes the Secondary without interrupting the Primary.

## **can\_sync**

If the `inconsistent` and `can_sync` flags are set, there is enough information in the Secondary SRL to make it consistent again. In this case, the RLINK does not need to be fully resynchronized and is still a suitable candidate for takeover.

## **cant\_sync**

If the RLINK flag is `cant_sync`, the RLINK is inconsistent and the Secondary needs a full synchronization before it can take part in replication again.

## **checkpoint**

A feature of VVR that synchronizes the Secondary using a block-level backup and restore method without interrupting the Primary.

## **checkpoint end**

Denotes the end point of the data for a Primary checkpoint. After a checkpoint attach of the RLINK, the Secondary becomes consistent when the checkpoint end is reached.



---

## **checkpoint start**

Denotes the starting point of the data for a Primary checkpoint. After the checkpoint attach of the RLINK, VVR starts sending writes performed after the checkpoint start.

## **CLEAN**

If the Primary RVG is in the CLEAN ACTIVE state, the RLINK must be attached.

## **copy-on-write**

A technique for preserving the original data. Before data is modified by a write operation, the original copy of data is copied to a different location.

## **consistent**

A flag indicating that data is recoverable by the system or application using it. For example, if the data contains a file system, the data is consistent if `fsck` can be run successfully on it. If the data contains a database, the data is consistent if the database recovery program can be run successfully and the database restarted.

## **Data Change Map (DCM)**

An object containing a bitmap that can be optionally associated with a data volume on the Primary RVG. The bits represent regions of data that are different between the Primary and the Secondary.

## **data volumes**

Volumes that are associated with an RVG and contain application data.

## **disaster recovery**

Disaster Recovery (DR) has a wide scope, which ranges from the storage of backup tapes off site to the setup and maintenance of a duplicate remote standby node. VVR aids disaster recovery by providing timely replication of data to a remote site.

## **distributed command**

A command or task that can be performed on an RDS from any host in the RDS environment; a related task is performed in sequence on all hosts in the RDS, if appropriate.

## **EMPTY**

If the Primary RVG state is EMPTY, an RLINK can be attached with no special options.



---

## FAIL

A Secondary RLINK can enter the FAIL state when the Secondary data volume has an error or when an ACTIVE Secondary RLINK is paused with the `-w` option. A Primary RLINK enters the FAIL state when the corresponding Secondary RLINK enters the FAIL state.

## failover

Failover is a term associated with the VERITAS Cluster Server environment. See [Primary takeover](#) for a description in the VVR environment.

## FastResync

A feature that is used to perform quick and efficient resynchronization of stale mirrors, and to increase the efficiency of the snapshot mechanism. Also see [Persistent FastResync](#) and [Non-Persistent FastResync](#).

## heartbeat protocol

The heartbeat protocol ensures that the nodes in an RDS will detect any network outage or a node crash. Once the Primary detects an outage, it begins a periodic attempt to reestablish a connection which continues until successful. Upon success, replication resumes automatically unless some interim administrative command or error prevents it.

## IBC

See [In-Band Control Messaging](#).

## In-Band Control Messaging

A facility that enables applications to inject control messages in the replication stream. The contents of the control message itself are application-defined and opaque to the replication process.

## inconsistent flag

A flag used for the following conditions:

- Inconsistent data on the Secondary, which indicates that the Secondary is not a viable takeover candidate.

- An RLINK in the FAIL state. Either the Secondary must be restored from a checkpoint or a full resynchronization must be performed.

This flag is set during DCM replay or a Primary checkpoint; it is cleared automatically when the DCM replay is complete or the checkpoint end is reached.



---

## latency high mark

Specifies the maximum number of waiting updates in the SRL before latency protection becomes active and writes stall or fail depending on the mode of the latency protection.

## latency low mark

Specifies the number of writes in the SRL when latency protection becomes inactive and writes proceed without waiting.

## latency protection

For RLINKs operating in asynchronous mode, which may fall behind, the latency protection attribute (`latencyprot`) of the RLINK is used to limit the maximum number of outstanding write requests. If latency protection is enabled, the maximum number of outstanding write requests cannot exceed the value set in `latency_high_mark`.

## latencyprot

See [latency protection](#).

## modes of replication

VVR replicates in *asynchronous* and *synchronous* modes. Each mode uses a different process to replicate data, and each deals differently with network conditions. See [asynchronous](#) and [synchronous](#) for more details.

## nmcom pool

The amount of buffer space available for updates coming in to the Secondary over the network.

## Non-Persistent FastResync

A form of [FastResync](#) that cannot preserve its maps across reboots of the system because it stores its change map in memory.

## object recovery

The process of making an object usable after a system crash.

## Passthru

Typically, writes to a Primary RVG are written to the SRL first, and then to the RLINKs and data volumes. If there is no SRL or the SRL is detached, writes are no longer written to the SRL and the RVG is in PASSTHRU mode. No replication takes place.

---

## Persistent FastResync

A form of [FastResync](#) that can preserve its maps across reboots of the system by storing its change map in a [DCO volume](#) on disk.

## plex

A copy of a volume and its data in the form of an ordered collection of subdisks. Each plex is a copy of the volume with which it is associated. The terms *mirror* and *plex* can be used synonymously.

## Primary checkpoint

A method for synchronizing a Secondary with the Primary without stopping the application. A command marks the start of the checkpoint. All Primary data volumes are backed-up and then the end of the checkpoint is marked. The data is restored on the Secondary and the Primary can then begin from the start of the checkpoint. The Secondary becomes consistent when the end of the checkpoint is reached.

## Primary pause

An administrator at the Primary volume node may pause updates to any particular RLINK of an RVG. During a pause, the Primary continues to keep a history of volume updates, but active update of the RLINK ceases and the network connection between the Primary and Secondary is broken. A Primary pause is intended as a maintenance feature and allows certain configuration changes to be made, such as a change to the network connecting two nodes.

## Primary migration

The Primary role of a volume can be migrated from a Primary node to a Secondary node, within certain restrictions. The process is manual and requires cooperative administrative action on the Primary and all Secondary nodes. During this process, update of the former Primary must be halted and cannot be resumed on the new Primary until the migration is complete.

The Primary role can only be migrated to an up-to-date RLINK that is consistent and is not in an error state. Any out-of-date Secondaries must be fully resynchronized with the new Primary.

## Primary node

The node on which the Primary RVG resides.

## Primary node crash

If a Primary node crash occurs, the Primary SRL and all data volumes in the RVG must be recovered. Both are recovered using VVR specific recovery, rather than the usual Volume Manager volume recovery.



---

## Primary node data volume failure

If there is an error accessing a Primary data volume, the data volume and the RVG are automatically detached, and the RVG state changes to `FAIL`. `RLINK`s are not affected. If the `SRL` volume was not empty at the time of the volume error, the updates continue to flow from the `SRL` to the Secondary `RLINK`s.

## Primary node SRL overflow

Because the Primary `SRL` is finite, prolonged halts in update activity to any `RLINK` can exceed the `SRL`'s ability to maintain all the necessary update history to bring an `RLINK` up-to-date. When this occurs, the `RLINK` is marked as `STALE` and requires manual recovery before replication can proceed.

## Primary Replicated Volume Group

In the VVR replication scheme, one copy of each replicated volume is designated the Primary volume, and the other copies are called Secondary volumes. Secondary volumes reside on different network hosts. Each Secondary volume group is represented on the Primary node by a single `RLINK` object, which serves as a proxy on the Primary node for the Secondary volume. Thus, when the Primary RVG needs to write data to a Secondary volume, it performs a write to the corresponding `RLINK` object. That object is then responsible for sending the request to the Secondary node. The `RLINK` object also provides an interface that allows for the sending of control messages to the remote RVG.

## Primary SRL failure

See [Passthru](#)

## Primary takeover

The Primary role can be arbitrarily taken over by a Secondary node. This process is similar to a Primary role migration but presumes that the old Primary is inoperable and unable to participate in the migration process.

Primary takeover is intended to support disaster recovery applications. Only a limited number of error scenarios prior to loss of the Primary node can prevent a takeover because they leave the `RLINK` in an inconsistent state. All such scenarios require the hardware failure of a data volume or `SRL`.

## RDS

See [Replicated Data Set](#)

## readback

Volume Replicator allocates memory for application writes from the `RVIO` pool and frees the memory when the write request is written to the Primary and all Secondary data volumes. When the memory pool space becomes low, VVR frees some memory space for



---

new write requests by postponing sending the data across asynchronous RLINKs. Later, these write requests are read back from the SRL when the RLINK is ready to send them. Such reads are called readbacks.

### **readback memory pool**

The amount of buffer space available for readbacks.

### **Replicated Data Set**

The group of the RVG on a Primary and its corresponding Secondary hosts.

### **Replicated Volume Group**

A component of VVR that is made up of a set of data volumes, one or more RLINKs, and an SRL.

### **Resynchronization**

The process of making the data on the Secondary identical to the data on the Primary. Resynchronization can be done without stopping the application by using a primary checkpoint or by using the Auto Synchronization feature.

### **RLINK**

RLINKs represent the communication link to the counterpart of the RVG on another node. At the Primary node a replicated volume object has one RLINK for each of its network mirrors. On the Secondary node a replicated volume has a single RLINK object that links it to its Primary.

### **RVG**

See [Replicated Volume Group](#)

### **RVIO memory pool**

The amount of buffer space allocated within the operating system to handle incoming writes.

### **Secondary checkpoint**

A method for backing up a Secondary to allow for quick recovery in the event of a data volume failure on the Secondary.

Allows volume-level backups of the RVG to be taken at the Secondary node



---

## Secondary data volume failure

Secondary data volume failure causes the RLINK to be put in the FAIL state. The Primary stops sending requests to the RLINK, but logging continues on the Primary node. When the failure has been corrected, it can be restored from a backup made earlier using a Secondary checkpoint.

## Secondary pause

An administrator at the Secondary node can pause updates to the RLINK. Unlike a Primary pause, the Primary-Secondary network connection is maintained. This enables the Secondary to notify the Primary when it wants updates of the RLINK to continue. A Secondary pause can be effected even when the Primary and Secondary have lost contact.

## Secondary node

The node to which the Primary data volumes are replicating.

## Secondary Replicated Volume Group

See [Replicated Volume Group](#).

## Secondary SRL volume failure

Secondary SRL volume failure is not a serious error because all the data is still available on the Primary. The failure will, however, cause the RLINK to be placed in a PAUSED state, just as if a system administrator had paused the volume.

## snapshot

A point-in-time image of a volume or file system that can be used as a backup.

## snapshot volume

An exact copy of a volume, at a specific point in time. The snapshot volume is created by breaking a snapshot plex from the original volume and placing it in a new volume, which is then used for online backup purposes.

## SRL overflow protection

For RLINKs with latency protection disabled, a final degree of protection is provided by SRL overflow protection (`srlprot`). When SRL overflow protection is enabled, the RLINK is protected from having so many outstanding write requests that it would overflow its SRL and require a full resynchronization.





---

## Storage Replicator Log (SRL)

Writes to the Primary RVG are saved in an SRL on the Primary side. The SRL is used to aid in error recovery, as well as to buffer writes when the system operates in asynchronous mode. Each write to a data volume in the RVG generates two write requests: one to the Secondary SRL and another to the Primary SRL.

## STALE

If an RLINK is STALE, the corresponding Secondary requires a full resynchronization to make it usable. RLINKS can enter the STALE state when they are detached manually by issuing the `vxrlink det` command or as a result of a Primary failure. In addition, when an RLINK is associated for the first time, its state changes from UNASSOC to STALE.

## synchronous

In synchronous mode, the Secondary is kept up-to-date with the Primary by waiting for each write request to reach the Secondary before the application sees the successful completion of the write on the Primary.

## unrecoverable errors

Some errors, in particular hard errors such as media failures, require manual intervention to repair. The chances of such failures can be minimized by using standard VxVM setups to maintain mirrored copies of each data volume and the SRL.

## update

Data from the Primary corresponding to the application writes sent to the Secondary for replicating the writes is referred to as an update.

## Volume Replicator Objects

The objects used for replication such as [Replicated Volume Group](#), [Storage Replicator Log \(SRL\)](#), [RLINK](#), and [Data Change Map \(DCM\)](#).

## write-order fidelity

A feature of VVR that applies writes on the Secondary in the same order as they are received on the Primary. This ensures that the data on the Secondary is consistent with the data on the Primary.





# Index

---

## A

- accessibility
  - assistive technology support, 415
  - dialog boxes, 412
  - keyboard navigation, 411
  - overview, 407
- acting Secondary, 232
- adding a data volume
  - using command line, 121
  - using VVR VEA, 274
- adding a Secondary
  - best practices, 46
  - using command line, 46
  - using VVR VEA, 258
- application-defined control messages, 384
- applications, on Secondary, 215
- assistive technology support, 415
- associating a data volume
  - using command line, 121
  - using VVR VEA, 274
- associating DCM to a data volume
  - using command line, 129
  - using VVR VEA, 274
- associating SRL to a Primary RVG
  - using command line, 44
  - using VVR VEA, 250
- asynchronous replication
  - data flow, 17
  - explained, 14
  - setting up, 50
- autodcm, srlprot setting, 29, 53
- autofb option, takeover command, 226
- automatic fast failback
  - using command line, 226
  - using VVR VEA, 289
- Automatic Synchronization
  - synchronizing Secondary example, 76
- automating firedrill, 184

- autosync flag, definition, 86

## B

- backup
  - block-level tape, 60, 66
  - primary checkpoint, 318
  - using snapshot and vradmin ibc, example, 213
  - using snapshot and vxibc, example, 381
- bandwidth limit, controlling
  - used for replication
    - using command line, 56
    - using VVR VEA, 256
  - used to synchronize volumes
    - using command line, 57
- Bandwidth Throttling
  - about, 55
  - disabling
    - using command line, 56
    - using VVR VEA, 256
  - See also* bandwidth limit
- bandwidth used by VVR, determining, 98
- block-level tape backup, 60, 66
- breaking off mirrors, 138

## C

- cache object, 136, 235
- can\_sync, 398
- cant\_sync, 398
- changing IP addresses, 145
- changing network ports for replication, 148
- changing replication settings, 143
- changing VVR tunables, 194
- checkdelete, vxrlink keyword, 156
- checkdelete, vxrvrg keyword, 155
- checkend, vxrvrg keyword, 328
- checkpoint attach
  - using VVR VEA, 262
- checkpoints



- backing up Secondary with, 191
- creating, 154
- data volume error
  - cleanup using, 318
  - recovery using, 318
- deleting, 155
- displaying list, 91
- ending, 155
- restoring Secondary
  - from Primary checkpoint, 192
  - from Secondary checkpoint, 191
- Secondary pause, 191
- understanding, 25
- viewing
  - using command line, 155
  - using VVR VEA, 273
- checkstart, vxrvrg keyword, 328
- commands
  - reference, 323
  - See also individual command names*
- comparing, snapshot methods, 40
- component volumes of a volume set,
  - associating to RDSs, 127
- configuration errors, recovery from, 306
- consistent data, concept, 10
- consistent RLINK, 193
- converting original Primary to Secondary
  - using command line, 217
  - using VVR VEA, 296
- cplist, vxrlink keyword, 91
- cplist, vxrvrg keyword, 91
- creating
  - cache object, 165
  - instant plex breakoff snapshot, 169
  - instant snapshot, 157
  - ISP cache object, 165
  - snapshot plexes, 168
  - snapshot volumes, 159
  - snapshots, overview, 157
  - space-optimized snapshot, 163
- creating a Primary RVG
  - using command line, 44
  - using VVR VEA, 250
- creating Replicated Data Set
  - using command line, 41
  - using VVR VEA, 249
- creating RVGs
  - using vradmin vs. vxmake, 43
- current data, concept, 10

## D

- Data Change Map, *see* DCM
- data differences, *see* data verification
- data flow
  - VVR asynchronous mode, 17
  - VVR synchronous mode, 19
- data loss, and takeover, 225
- data on Secondary, using for off-host
  - processing, 199
- data verification
  - concept, 185
  - performing, 185
- data volume errors
  - missing, 309
  - name mismatch, 311
  - resize, 310
- data volumes
  - administering
    - using command line, 121
    - using VVR VEA, 274
  - associating
    - using command line, 121
    - using VVR VEA, 274
  - dissociating or removing
    - using command line, 134
    - using VVR VEA, 276
  - mapping Secondary name, 401
  - resizing
    - using command line, 132
    - using VVR VEA, 275
  - verifying data
    - about, 185
    - before adding volumes to RDS, 123
    - using offline verification, 187
    - using online verification, 185
- data, consistent vs. up-to-date, 10
- DCM
  - and fast failback, 226
  - associating to a data volume, 129
  - concept, 9
  - required for takeover, 225
- DCM playback
  - after SRL overflow, 135
  - during failback, 226, 232
- DCM size, 130
- dcm, srlprot setting
  - about, 29
  - changing, 53
- dcm\_logging flag, definition, 136



---

- Decision Support, example, 211, 378
- detailed view, VEA, 247
- difference-based synchronization
  - about, 70
  - failing back using, 237
- disabling Bandwidth Throttling, 56
- disaster recovery
  - failed Primary, 223
  - verifying readiness, 33, 183
- Disk Group Split and Join, synchronizing Secondary using, 68
- disk groups
  - displaying disk group ID, 48
- displaying
  - checkpoints, list, 91
  - data volumes in RDS, 84
  - RVG and RDS information, 83
- dissociating data volumes
  - using command line, 134
  - using VVR VEA, 276
- DR readiness, verifying, 183
- DRL, DCM compared to, 9
- DSS, *See* Decision Support

## E

- error messages, 352
- errors
  - configuration, 306
  - data volume size mismatch, 310
  - during an initial RLINK attach, 306
  - during modification of an RVG, 309
  - missing data volume, 309
  - name mismatch, data volume, 311
  - vradmind daemon, 354
- event notification
  - using command line, 117
  - using VVR VEA
    - configuring SMTP mail server, 300
    - explained, 299
    - recipients, setting up, 300
    - rule manager, 299
    - SNMP MIB files for VVR, 302
- event types, 117
- examples
  - using command line
    - creating an RDS, 75
    - migrating Primary, 219
    - off-host processing, 211, 378
    - setting up replication, 76

- synchronizing the Secondary, 67
  - taking over from failed Primary, 228, 229
- using VVR VEA
  - failing back using automatic synchronization, 297
  - failing back using fast failback, 294
  - setting up replication, 251

## F

- fail
  - latencyprot setting, 31
  - mode of replication, 27
  - srprot setting, 53
- FAIL state, 397, 399
- failback logging
  - about, 226
  - example, 234
- Failback options, VVR VEA
  - Auto fast failback, 289
  - Fast failback, 289
  - No fast failback, 290
- failback, VVR VEA, 296, 297
- failed Primary, example takeover, 228, 229
- failing back
  - using command line
    - about, 231
    - fast failback versus difference-based synchronization, 231
    - using difference-based synchronization, 237
    - using fast failback or failback logging, 232
  - using VVR VEA
    - about, 292
    - using automatic synchronization, 295, 297
    - using fast-failback synchronization, 293
- fast failback feature
  - about, 226
  - enabling with vradmind takeover, 226
  - versus difference-based synchronization, 231
- Fast failback option, VVR VEA, 289
- FastResync, 182
- features of VVR, 1
- firedrill, performing, 184
- firewalls



- using VVR in, 47
- VVR ports in, 149
- flag settings
  - RLINK, 86
  - RVG, 89
- flags and definitions
  - autosync, 86
  - can\_sync, 398
  - cant\_sync, 398
  - Data Change Map, 136
  - inconsistent, 398
  - resync\_active, 137
  - resync\_paused, 87, 137
  - resync\_started, 87
  - SRL Overflow Protection With DCM, 136
- freezing replication
  - about, 199
  - before taking snapshot, 160
  - off-host processing, 202, 205
  - using ibc messaging, 372
- full synchronization
  - about, 64
  - example, 66

## G

- global mapping, 401
- Graphical User Interfaces *See* GUIs, VVR
- GUI, VVR
  - about, 41
- GUIs, VVR
  - VEA, 241

## H

- heartbeats, ports used for, 149
- high mark, latency, setting
  - using command line, 51
  - using VVR VEA, 256
- host names, changing, 145

## I

- IBC Messaging
  - API, 384
  - explanation, 201
  - function, 384
  - off-host processing, using for, 201
  - overview, 199, 372
  - vradmin ibc, using for off-host processing, 206
  - vxibc, individual commands, 373

- inactive RLINKs, 396
- In-Band Control Messaging, *see* IBC Messaging
- inconsistency
  - inconsistent flag, 398
  - inconsistent RLINKs, 398
  - RLINK during restores, 193
- initial VVR setup
  - using command line, 73
  - using VVR VEA, 249
- interfaces of VVR, about, 41
- IP addresses, changing, 145
- ISP cache object
  - creating, 165

## K

- kernel states (KSTATEs), 393

## L

- latency high mark, defined, 420
- latency low and high mark, setting
  - using command line, 51
  - using VVR VEA, 256
- latency low mark, defined, 420
- latency protection
  - about, 31
  - disabling
    - using command line, 52
    - using VVR VEA, 256
  - enabling
    - using command line, 51
    - using VVR VEA, 256
- local host, definition, 12
- local RVG, represents RDS, 12
- log overflow, STALE, 396
- low mark, latency, setting
  - using command line, 51
  - using VVR VEA, 256

## M

- make Secondary
  - using command line, 237
  - using VVR VEA, 296
- memory statistics, displaying, 332
- messages
  - debugging, 346
  - kernel, 333
  - utility error, 346
- MIB files, VVR VEA, 302
- migrate, vradmin keyword, 217



- migrating Primary
  - example, 219
  - overview, 216
  - using command line, 217
  - using VVR VEA, 287
- mirrors, breaking off before
  - synchronization, 138
- missing data volume error, 309
- modes of replication, setting, 50
- most up-to-date Secondary, identifying, 114
- multiple Secondaries in RDS, migrating
  - Primary with, 218

## N

- network bandwidth used by VVR
  - determining, 98
- network performance data, 94, 97, 111
- network ports, changing, 148
- network transport protocol
  - attribute, TCP, 54
  - attribute, UDP, 54
  - planning, 47
  - setting, 54
- No fast fallback option, VVR VEA, 290
- notify, events, 117

## O

- object states, VVR, 393
- object tree, VEA, 246
- off
  - latencyprot setting, 31
  - mode of replication, 27, 50
  - srprot setting, 53
- off-host processing
  - backup example, 381
  - Decision Support example, 211, 378
  - examples, 211, 378
  - IBC Messaging, using, 201
  - tasks to perform, 205
  - trial failover example, 382
- offline data verification
  - concept, 185
  - using command line, 187
- onfreeze script
  - about, 205
  - creating, 207
- onfreeze task, off-host processing, 205
- online backup, 193
- online data validation, 185
- overflow protection, SRL, 53

- override
  - latencyprot setting, 31
  - mode of replication, 27, 50
  - srprot setting, 29, 53

## P

- packet size, setting, 54
- pause, vxlink keyword, 397
- pause/resume, temporary states, 399
- pauserep, vradm keyword, 143
- pausing replication
  - using command line, 143
  - using VVR VEA, 279
- Pausing state, 399
- performing
  - offline data validation, 185
  - online data validation, 185
- performing firedrill, 184
- ports used by VVR, changing, 148
- Primary checkpoint
  - cleanup using, 318
  - deleting, 155
  - restoring from, 192
  - restoring Secondary from, 397
- Primary host
  - concept, 11
  - failure recovery, 312
- Primary role takeover
  - example, 228, 229
  - overview, 223
- Primary RVG
  - creating
    - using command line, 44
    - using VVR VEA, 258
  - removing
    - using command line, 153
    - using VVR VEA, 285
- Primary SRL
  - error, cleanup after, 314
  - header error, 316
  - overflow recovery, 316
  - primary\_datavol field, 401
  - printrvg, vradm keyword, 83
  - printvol, vradm keyword, 84
  - protocol, network, 47

## R

- RAID-5 volumes and VVR, 44
- RDS
  - administering data volumes



---

- using command line, 121
  - using VVR VEA, 274
- associating a volume set, 127
- associating volumes
  - using command line, 121
  - using VVR VEA, 274
- concept, 11
- creating
  - using command line, 44
  - using VVR VEA, 258
- displaying data volumes in RDS
  - using command line, 84
  - using VVR VEA, 270
- displaying information
  - using command line, 83
  - using VVR VEA, 265
- displaying replication status
  - using command line, 103
  - using VVR VEA, 265
- removing
  - using command line, 153
  - using VVR VEA, 285
- removing or dissociating data volumes
  - using command line, 134
  - using VVR VEA, 276
- removing Primary from
  - using command line, 153
  - using VVR VEA, 285
- resizing data volumes
  - using command line, 132
  - using VVR VEA, 275
- resizing SRL
  - using command line, 141
  - using VVR VEA, 277
- synchronizing RVG
  - using VVR VEA, 297
- RDS View, VVR VEA, 267
- recipients, setting up using VVR VEA, 300
- RECOVER state, 399
- recovering from failed Primary, 223
- recovering RLINK, 330
- Recovery Point Objective (RPO), 225
- Recovery Time Objective (RTO), 225
- region size, 130
- remote administration using VEA, 242
- removing data volume
  - using command line, 134
  - using VVR VEA, 276
- removing Primary
  - using command line, 153
  - using VVR VEA, 285
- removing Secondary
  - using command line, 152
  - using VVR VEA, 284
- replacement for Primary
  - finding most suitable, 225
- Replicated Data Set, *see* RDS
- replication
  - freezing, 160, 202
  - pausing, 143
  - resuming, 143
  - resuming after migration, 220, 222
  - starting, 58
  - status, displaying, 103
  - stopping to Secondary, 144
  - terms defined, 10
  - unfreezing, 163, 375
- replication settings
  - bandwidth limit, 55
  - defaults, 49
  - latency protection, 51
  - overview, 143
  - packet size, 54
  - replication mode, 50
  - SRL overflow protection, 53
- repstatus, vradm keyword, 103
- Resize SRL task, VVR VEA, 277
- Resize Volume task, VVR VEA, 275
- resizevol, vradm keyword, 132
- resizing a file system, 132
- resizing data volumes
  - about, 132
  - using command line, 133
  - using VVR VEA, 275
- resizing SRL
  - using command line, 141
  - using VVR VEA, 277
- restore
  - restore command, 399
  - restore, vxrlink command, 192, 193
  - RESTORING state, 399
  - Secondary RLINK, 193
- restore, vxrlink keyword, 330, 397
- resume
  - resume command, 399
  - RESUMING state, 399
- resume, vxrlink keyword, 330
- resumerep, vradm keyword, 143





- 
- resuming replication
    - after migration, 220, 222
    - using command line, 143
  - resync
    - resync\_active flag, definition, 137
    - resync\_paused flag, definition, 87, 137
  - resync\_started flag, definition, 87
  - resyncfromreplica
    - recovering failed Secondary data
      - volumes, 181
    - recovering from logical corruption of data, 180
    - using snapback, 180
  - right pane, VVR VEA, 247
  - RLINK
    - concept, 8
    - configuring, 49
    - creating, 332
    - displaying status, 85
    - dissociating from RVG, 330
    - flag settings, 86
    - inactive, 396
    - STALE, 316
    - states, 395
    - status command, 78
  - Rule Manager, VVR VEA, 299
  - RVG
    - associating to clear PASSTHRU, 314
    - concept, 6
    - concept, VVR VEA, 243
    - displaying status, 88
    - flag settings, 89
    - start, 328
- S**
- Secondary
    - See also* synchronizing Secondary
    - adding, 46
    - backing up, 191
    - removing from RDS, 152
    - using data, 199
  - Secondary checkpoint
    - backing up the Secondary, 191
    - deleting, 156
    - recovery using, 318
  - Secondary host, concept, 11
  - Secondary SRL
    - header error, 320
    - use of, 13
    - volume error, 319
  - Secondary, acting, 232
  - Secondary, identifying most up-to-date, 114
  - Secondary, removing, 152
  - Secondary, transferring applications to, 215
  - secondary\_log\_err flag, 319
  - secondary\_paused flag, 319
  - set, vradm keyword, 49
  - setting bandwidth limit
    - using command line, 56
    - using VVR VEA, 256
  - setting network transport protocol, 54
  - setting packet size, 54
  - size of DCM, 130
  - SMTP mail server, 300
  - snapback operation, about, 179
  - snapback, vxrv keyword, 179, 329
  - snapprnt, vxrv keyword, 329
  - snaprefresh, vxrv keyword, 329
  - snaprestore, vxrv keyword, 329
  - snapshot file systems, 424
  - snapshot methods
    - overview
      - instant full, 35
      - instant plex breakoff, 36
      - instant space-optimized, 35
  - snapshot plexes, creating, 177
  - snapshot, vxrv keyword, 330
  - snapshots
    - before DCM replay, 232
    - before takeover, 225
    - comparing methods, 40
    - overview, 157
    - traditional method, 34
    - understanding, 33
    - using with fast failback, 235
    - using with resynchronization, 136
  - snapshots of RVG
    - creating
      - cache object, 165
      - instant plex breakoff, 169
      - instant snapshot, 157
      - snapshot plexes, 168
      - snapshot volumes, 159
      - space-optimized, 163
    - destroying, 176
    - displaying information, 175
    - online data validation, 185
    - reattaching, 179



- refreshing, 171
  - restoring, 173
  - resuming replication, 163, 167, 170
  - taking, 178
  - unfreezing replication, 163
  - SNMP MIB files, VVR VEA, 302
  - splitter, VVR VEA, 247
  - SRL
    - associating to a Primary RVG
      - using command line, 44
      - using VVR VEA, 250
    - concept, 7
    - resizing
      - using command line, 141
      - using VVR VEA, 277
  - SRL header error, 316
  - SRL overflow protection
    - disable, 30
    - explanation, 53
    - modes of, 29, 53
    - preventing, 316
    - with DCM, flags and definition, 136
  - SRL playback, 232
  - srlprot
    - srlprot=dcm, override, autodcm, 53
  - STALE RLINK, 396
  - starting replication, 58
  - starting replication using VVR VEA
    - about, 260
    - options, 261
  - startrep, vradm keyword, 62
  - statistics, displaying for VVR, 92
  - stats, vxlink keyword, 111, 331
  - stats, vxrv keyword, 329
  - status, vxlink keyword, 78, 331
  - Stop Replication task, VVR VEA, 280
  - stopping replication to Secondary
    - using command line, 144
    - using VVR VEA, 280
  - stoprep, vradm keyword, 144
  - Synchronize RVG task, VVR VEA, 297
  - synchronizing Secondary
    - automatic synchronization
      - example, 76
      - using VVR VEA, 261
    - backup and checkpoint
      - using VVR VEA, 262
    - difference-based synchronization
      - about, 70
      - example, 71
      - using command line, 70
    - full synchronization
      - about, 64
      - example, 66
      - using command line, 65
    - incrementally after SRL overflow
      - about, 135
      - using command line, 137
      - using VVR VEA, 285
  - synchronizing volumes
    - about, 124
    - using difference-based synchronization, 125
    - using full synchronization, 125
  - synchronous mode, data flow, 19
  - synchronous replication
    - explanation, 14
    - setting up, 51
    - synchronous=off, 50
    - synchronous=override, 51
  - syncrv, vradm keyword
    - and volume sets, 127
    - difference-based synchronization, 70
    - full synchronization, 64, 65
  - syncvol, vradm keyword, 124, 127
- T**
- takeover, *see* taking over from failed
    - Primary
  - takeover, vradm keyword, 224
  - taking over from failed Primary
    - example, 228, 229
    - overview, 223
    - using command line, 224
    - using VVR VEA, 289
  - tape backups
    - block-level, 60, 66
  - TCP protocol, 47
  - Technical assistance, xv
  - terms
    - definitions, 11
    - replication, in VVR context, 10
  - traditional snapshot method, 34
  - transferring Primary role
    - failing back, 231
    - methods, 215
    - migrating, 216
    - taking over, 223

---

- tree, VEA, 246
- trial failover, example, 382
- tunables, VVR
  - changing
    - using vxtune, 194
  - displaying, 97

## U

- UDP protocol, 47
- unfreezing replication
  - after snapshot, 163
  - using vxibc unfreeze, 375
- unrecoverable I/O error, 397
- update ID, displaying, 114
- updates, vxrlink keyword, 114, 331
- up-to-date data, concept, 10
- up-to-date Secondary, determining, 114
- usage types, VVR volumes, 44
- utility error messages, 352
- utility, vrnotify, 117
- utility, vxrlink stats, 111
- utility, vxtune, 195

## V

- verify, vradadmin keyword, 123
- verify, vxrlink keyword, 331
- verifydata, vradadmin keyword, 185
- verifying DR readiness
  - Failover, 183
  - Firedrill, 184
  - offline data validation, 187
  - online data validation, 185
  - overview, 183
- verifying secondary data, 185
- verifying volumes
  - about, 185
  - before adding to RDS, 123
  - using offline data verification, 187
  - using online data verification, 185
- VERITAS Volume Replicator,
  - accessibility, 407
- viewing
  - consolidated statistics, 93
  - data volume statistics, 95
  - data volumes in RDS, 84
  - flag settings, 86
  - memory tunable parameters,
    - statistics, 97
  - network bandwidth used by VVR, 98
  - replication status, 85

- RLINK statistics, 94
- RVG and RDS information, 83
- SRL volume statistics, 96
- status of individual RVG, 88
- status of specific RLINK, 85
- VVR statistics, 92
- vol\_max\_nmpool\_sz tunable, 112
- Volume Manager, VVR as option, 1
- volume set, associating to an RDS, 127
- volumes
  - associating to an RDS
    - using command line, 121
    - using VVR VEA, 274
  - removing or dissociating from an RDS
    - using command line, 134
    - using VVR VEA, 276
  - resizing
    - about, 132
    - using command line, 133
    - using VVR VEA, 275
- volumes, synchronizing
  - about, 124
  - using difference-based synchronization, 125
  - using full synchronization, 125
- volumes, verifying
  - before adding to RDS, 123
  - using offline data verification, 187
  - using online data verification, 185
- vradadmin command
  - command reference, 323
  - error messages, 354
- vradadmin utility
  - addsec, 46
  - addvol, 121
  - changeip, 145
  - createpri, 44
  - creating RVGs, 43
  - delpri, 153
  - delsec, 152
  - delvol, 134
  - fbsync, 226, 232, 235
  - ibc
    - compared with vxibc, 371
    - using off-host processing, 199
  - makesec, 237
  - migrate, 217
  - pauserep, 143
  - printrvg, 83, 358



---

- printvol, 84
- repstatus, 103
- resizesrl keyword, 141
- resizevol, 132
- resumerep, 143
- resync keyword, 135
- set, 49
- set keyword, 143
- startrep, 62
- stoprep, 144
- syncrvg
  - and volume sets, 127
  - difference-based synchronization, 70
  - full synchronization, 64, 65
- syncvol, 123, 124, 127
- takeover, 224
- verify, 123
- verifydata, 185
- vradmind, display commands, 83
- vradmind daemon
  - errors, 354
  - restarting, 92
- vrnotify utility, 117
- vrport command, 149
- vrstat command, 92
- VVR GUIs
  - VEA, 241
- VVR terms, defined, 11
- VVR tunables
  - changing, 194
  - displaying, 97
- VVR VEA
  - accessibility
    - assistive technology support, 415
    - dialog boxes, 412
    - keyboard navigation, 411
  - bandwidth limit, determining, 273
  - detailed view, 247
  - example scenario, 251
  - failback options, about, 289
  - right pane, 247
  - using to control bandwidth limit, 256
  - windows, resizing, 247
- VVR VEA tasks
  - Associate Volume, 274
  - Dissociate Volume, 276
  - Make Secondary, 296
  - Migrate, 287
  - Pause Replication, 279
  - Remove Primary, 285
  - Remove Secondary, 284
  - Resize SRL, 277
  - Start Replication, 260
  - Stop Replication, 280
  - Synchronize RVG, 297
  - Takeover
- VVR VEA wizards
  - Create Replicated Data Set, 249
  - Resize Volume, 275
- VVR, features, 1
- VVR, option of VxVM, 1
- vxassist command
  - addlog keyword, 129
  - growto keyword, 142
  - make keyword
    - associating DCM, 129
    - creating cache volumes, 165
    - creating snapshot volumes, 162
  - maxgrow keyword, 353
  - mirror keyword, 139
  - snapback keyword, 180
  - snapshot keyword, 177
  - snapstart keyword, 177
- vxconfigd, daemon
  - messages, 352
- vxdg command
  - free keyword, 141
  - split keyword, 68, 79
  - upgrade keyword, 159
- vxedit command
  - rename keyword
    - for plex-breakoff snapshots, 170
    - in Disk Split and Join, 68
    - renaming volumes during RLINK recovery, 308
  - rm keyword, 139
  - set keyword
    - mapping volume names with, 304
    - primary\_datavol, 402
    - setting synchronous=fail, 50
    - set remote, 403
- vxibc utility
  - about, 371
  - send command, 374
  - unfreeze command, 375
- vxmake command
  - creating cache objects
    - for space-optimized snapshots, 165

- setting attributes, 163
  - creating RLINK, 332
  - creating RVGs, 43
  - creating volumes during plex recovery, 139
  - rlink keyword, 403
- vxmemstat command
  - command reference, 332
  - versus vrstat command, 92
- vxnetd, starting, 341
- vxplex command, 138
- vxprint command
  - displaying primary\_datavol field, 402
  - displaying VVR status, 85
- vxrlink att, 393
- vxrlink command
  - assoc keyword, 330
  - att keyword, 142, 330
  - checkdelete keyword, 156
  - cplist keyword, 91
  - det keyword, 139, 142, 330
  - dis keyword, 330
  - pause keyword
    - creating Secondary checkpoint, 154
    - RLINK FAIL state, 397
    - using Secondary checkpoints, 192
  - recover keyword
    - command reference, 330
    - two-phase transactions in, 399
  - restore keyword, 192, 193, 330, 397
  - resume keyword, 330
    - ending Secondary checkpoint, 155
    - using Secondary checkpoints, 192
  - stats keyword
    - displaying network statistics, 111
    - reference, 331
    - versus vrstat, 92
  - status keyword
    - displaying status of RLINK, 89
    - reference, 331
    - synchronization after SRL overflow, 137
  - updates keyword, 114, 331
  - verify keyword, 331
  - versus vrstat, 92
- vxrlink recover, 394
- vxrsyncd daemon, ports, 148
- vxrvg command
  - checkdelete keyword, 155
  - checkend keyword, 155
  - checkstart keyword, 154
  - command reference, 328
  - cplist keyword, 91
  - getdatavols keyword, 89
  - recover keyword, 399
  - reference, 328, 329
  - snapback keyword
    - reattaching instant snapshot plexes, 171
    - reattaching traditional snapshot plexes, 179
    - using FastResync, 182
  - snapdestroy keyword, 176
  - snapprint keyword, 175
  - snaprefresh keyword, 171
  - snaprestore keyword, 173
  - snapshot keyword
    - breaking off mirrors, 138
    - creating instant snapshots, 157
    - creating traditional snapshots, 177
    - using FastResync, 182
    - using with IBC, 380
  - start keyword, 393
  - stop keyword, 174
- vxsnap command, 162
- vxstat command, versus vrstat command, 92
- vxtune utility, 195
- VxVM, VVR as option, 1
- vxvol command
  - aslog keyword
    - associating the SRL, 142, 314
  - assoc keyword
    - associating a new volume, 308
    - during disaster recovery, 139
    - setting Primary volume name on Secondary volume, 402
  - dis keyword
    - dissociating the SRL, 142, 314
    - during disaster recovery, 139
  - set keyword
    - enabling FastResync, 182
  - start keyword
    - during disaster recovery, 139
    - starting data volume, 313
    - starting SRL, 314
- vxvoladm command, 165
- vxvset command, 127



---

**W**

write-order fidelity, concept, 10

writing to Secondary data volumes, 215

