

Veritas™ Volume Replicator Advisor User's Guide

AIX, HP-UX, Linux, Solaris

5.0

Veritas Volume Replicator Advisor User's Guide

Copyright © 2006 Symantec Corporation. All rights reserved.

Veritas Volume Replicator 5.0

Symantec, the Symantec logo, Veritas, and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

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

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

The Licensed Software and Documentation are deemed to be "commercial computer software" and "commercial computer software documentation" as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014
www.symantec.com

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

AIX is a registered trademark of IBM Corporation.

HP-UX is a registered trademark of Hewlett-Packard Development Company, L.P.

Linux is a registered trademark of Linus Torvalds.

Solaris is a trademark of Sun Microsystems, Inc.

Licensing and registration

Veritas Volume Replicator is a licensed product. See the *Veritas Volume Replicator Installation Guide* for license installation instructions.

Technical support

For technical assistance, visit <http://support.veritas.com>. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Chapter 1	Introducing Volume Replicator Advisor (VRAdvisor)	
	Audience	7
	Related Veritas documents	8
	Overview of VRAdvisor	8
	How VRAdvisor works	9
	Data collection	9
	Data analysis	9
	What-if analysis	10
Chapter 2	Installing Volume Replicator Advisor (VRAdvisor)	
	Installing VRAdvisor on Solaris	11
	Uninstalling VRAdvisor on Solaris	11
	Installing VRAdvisor on Windows	12
	Uninstalling VRAdvisor	13
Chapter 3	Collecting the sample of data	
	Collecting the sample of data on UNIX	16
	Supported locales	16
	Collecting data using the VRAdvisor wizard	17
	Collecting data using the vxstat command	18
	Collecting data using the data collection script	19
	Collecting the sample of data on Windows	22
	Collecting the sample of data using the VRAdvisor wizard	22
	Collecting the sample of data using the diskStats command	25
Chapter 4	Analyzing the sample of data	
	Analyzing the collected data	28
	Understanding the results of the analysis	33
	Viewing the analysis results	33
	Recalculating the analysis results	34
	Recording and viewing the results	38

Appendix A Sizing the SRL

Overview	39
Peak usage constraint	40
Synchronization period constraint	42
Secondary backup constraint	43
Secondary downtime constraint	44
Additional factors	44
Example	46

Index	49
-------------	----

Introducing Volume Replicator Advisor (VRAdvisor)

Veritas Volume Replicator Advisor (VRAdvisor) is a planning tool that helps you determine an optimum Veritas Volume Replicator (VVR) configuration. This document assumes that you understand the concepts of VVR.

This document provides information on installing and using this tool on different platforms. Wherever applicable, the information that is specific to a platform has been appropriately indicated. For Windows, note that the Veritas Volume Manager (VxVM) has been renamed to Veritas Storage Foundation for Windows (VSFW) from Release 4.1 onwards.

For a description of the concepts of VVR, refer to the *Veritas Volume Replicator Administrator's Guide*. For information about planning and tuning VVR, refer to the *Veritas Volume Replicator Planning and Tuning Guide*.

Audience

This guide is intended for system administrators who are responsible for installing, configuring, and setting up replication using VVR. This guide assumes that the user has:

- A basic understanding of system administration.
- Working knowledge of the VVR product.

This document guides you through the process of installing VRAdvisor and then evaluating various parameters using the data collection and data analysis process. This document describes procedures using both the graphical user

interface and the command-line interface, as applicable, on the different platforms.

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.

Overview of VRAdvisor

Planning is the key to successfully configuring VVR. To set up an optimum configuration, you must understand the components of VVR and their interactions with each other. In addition, you must consider the factors that are specific to your environment while planning your VVR configuration. The important factors to consider while planning include:

- Needs and constraints of the business
- Application characteristics
- Mode of replication
- Network characteristics

These factors are dependent on each other and these dependencies must be considered during planning. For example, if your business requires that the data on the Secondary to be as up to date with the Primary as possible, you must choose synchronous mode and provide enough network bandwidth to handle the peak application write rate on the Primary. Or, if the available network bandwidth is less than the peak write rate of the application, you must choose asynchronous mode of replication. Also, the size of the Storage Replicator Log (SRL) must be able to handle the Secondary outages and network outages for the given application characteristics. VRAdvisor considers these dependencies and enables you to determine the parameters to suit your VVR environment.

VRAdvisor does the following:

- ✓ Collects a sample of data that reflects the application characteristics.
- ✓ Analyzes the sample of the application characteristic and calculates the size of the SRL and the network bandwidth required for replication.
- ✓ Enables you to perform a What-if Analysis by varying the needs and constraints of your business, based on your future requirements.

Note that the replication log of VVR is referred to as SRL (Storage Replicator Log) on UNIX and as Replicator Log on Windows. The terms SRL and Replicator Log are used interchangeably in the document.

How VRAdvisor works

Using VRAdvisor for planning involves collecting a sample of data that represents the application write rate and analyzing this sample of data based on factors, such as the network bandwidth and network outage. VRAdvisor considers the worst case situations when analyzing data, which results in an optimum configuration for VVR. Working with VRAdvisor involves:

- ✓ [Data collection](#)
- ✓ [Data analysis](#)
- ✓ [What-if analysis](#)

Data collection

VRAdvisor uses a sample of data for analysis; the sample of data must be available in an appropriate format required by VRAdvisor. To collect a sample of data that represent the application write rate, we recommend that you collect the sample of data for a period of seven to fourteen days. Make sure that the collection period includes times of peak usage for your application, so that the collected data reflects your environment.

In the data collection mode, VRAdvisor collects the sample of data in the appropriate format required by VRAdvisor. You can also collect the sample of data using the data collection script provided. The data collection script uses the appropriate command at the operating system level to collect the data, and also converts the data to the appropriate format required by VRAdvisor. For more information, see “[Collecting the sample of data](#)” on page 15.

Data analysis

In the data analysis mode, VRAdvisor analyzes the sample of data that you have collected, based on the following factors specified by you:

- Available network bandwidth
- Network outage duration
- Secondary outage duration

After analyzing the data, VRAdvisor displays a graphical as well as textual representation of the results in a separate window. For more information, see “[Analyzing the sample of data](#)” on page 27.

What-if analysis

The What-if analysis feature enables you to perform additional calculations, to plan for future requirements or alternative scenarios. You can vary the parameters and recalculate the results according to different criteria. For example, you can vary the network bandwidth parameter to see what effect it would have on the SRL size. Or, you can specify a potential SRL size and see how much network bandwidth would be required for that SRL size. For more information, see [“Performing What-if analysis”](#) on page 35.

Installing Volume Replicator Advisor (VRAdvisor)

This chapter explains how to install Veritas Volume Replicator Advisor on the Solaris and Windows operating systems.

VRAdvisor is supported on the following operating systems:

- Solaris 8, Solaris 9, and Solaris 10
- Windows XP
- Windows 2000 Professional, Windows 2000 Server, Windows 2000 Advanced Server, Windows 2000 Datacenter Server

Note: If Windows 2000 is installed on your system, you must have Service Pack (SP) 4 installed on it.

- Windows Server 2003 (32-bit): Standard Edition, Web Edition, Enterprise Edition, Datacenter Edition

Installing VRAdvisor on Solaris

To install, run the following command:

```
# pkgadd -d volume_replicator/tools/vradvisor/pkgs VRTSvradv  
where VRTSvradv is the package name.
```

Uninstalling VRAdvisor on Solaris

To uninstall VRTSvradv, run the following command:

```
# pkgrm VRTSvradv
```

Installing VRAdvisor on Windows

This section gives instructions on installing `VRAdvisor` on Windows. The procedure to install on Windows 2000 and Windows Server 2003 is the same, except that on a Windows 2000 system, a reboot is required after the installation completes.

Note: VRAdvisor is not installed as a part of the common installation process that uses the product installer. To install Veritas Volume Replicator Advisor, follow the procedure in this section.

Note: Although VRAdvisor is supported in a non-English locale, the wizards are still displayed in English.

To install VRAdvisor

- 1 If a previous version of VRAdvisor is installed, remove the existing VRAdvisor before installing VRAdvisor.
- 2 Navigate to the `windows` directory on the CD.
- 3 Run the `VRTSvradv.msi` from the `windows` directory.
The installation wizard is launched. A message indicates that the `VRAdvisor` setup file is checking for the necessary parameters before starting the installation process.
- 4 In the Welcome page, click **Next**.
- 5 In the Customer Information page, enter your user name and organization, and click **Next**.
- 6 Complete the Destination Folder page as follows:
 - To install `VRAdvisor` in the default directory `c:\Program Files\VERITAS\Volume Replicator Advisor`, click **Next**.OR
 - a To choose another location for installing `VRAdvisor`, click **Change**.
 - b In the Change Current Destination Folder page, in the **Folder name** field, enter the complete path to the directory where you want the `VRAdvisor` package to be installed. You can also use the browse button to navigate to the required directory. Click **OK**.
 - c In the Destination Folder page, click **Next**.

- 7 In the Ready to Install the Program page, click **Install** to proceed with the installation.
The Installing VERITAS Volume Replicator Advisor page appears. This page displays a progress bar to indicate that the installation is in progress.
After the installation completes, a message indicates that the installation was successful.
- 8 Click **Finish**.
- 9 If required, a message prompts you to restart the system. Click **Yes** to reboot the system now. Click **No** to reboot it later.
On Windows 2000, installing VRAdvisor requires a reboot to activate disk performance counters for logical volumes. By default, Windows 2000 activates only the Physical Disk performance counters. To collect sample data, VRAdvisor requires the Logical Disk performance counters. The VRAdvisor installer runs the `diskperf -yv` command to activate the Logical Disk performance counters; however, the system must be restarted before the counters are enabled.
On Windows 2003 or Windows XP, a reboot is not required to enable disk performance counters.

Uninstalling VRAdvisor

- 1 To uninstall VRAdvisor, select **Settings > Control Panel** from the Windows **Start** menu.
- 2 Select **Add or Remove Programs**.
- 3 Select **VERITAS Volume Replicator Advisor** from the list of programs.
- 4 Click **Remove**. Windows prompts you to confirm that you want to remove Veritas Volume Replicator Advisor.
- 5 Click **Yes**. The VERITAS Volume Replicator Advisor dialog box appears. The progress bar on the VERITAS Volume Replicator Advisor dialog box indicates that the removal is in progress.
After the uninstallation procedure completes, the Add or Remove Programs dialog indicates that the Veritas Volume Replicator Advisor program has been removed successfully.

Collecting the sample of data

You need to collect data write samples that can be used with the VRAdvisor wizard. VRAdvisor uses the sample of data for analysis.

Best practices

- We recommend that you collect the samples of data using the volumes that are to be part of the VVR configuration you are planning to set up.
- To collect a representative sample of data, we recommend that you collect the sample of data over a period of 7 to 14 days.

Note: The data must be collected for a minimum of seven days.

- Make sure that the collection period includes times of peak usage for your application, so that the collected data reflects your actual requirements.
- In a shared disk group, collect a sample of data on each node. Start the collection at the same time on each node, for the same period. The resulting files can then be analyzed together.

VRAdvisor calculates an optimum size of the Storage Replicator Log (SRL) and the network for your VVR configuration using a sample of the write statistics.

Depending on the operating system on which you are collecting data, you can either collect the sample of data using the VRAdvisor wizard, commands, or the data collection script. For details, refer to the section for your platform.

- [“Collecting the sample of data on UNIX”](#) on page 16
- [“Collecting the sample of data on Windows”](#) on page 22

Collecting the sample of data on UNIX

VRAdvisor can be used to collect a sample of data for analysis. You can collect data using the VRAdvisor wizard, vxstat, or the data collection script. To use the VRAdvisor wizard to collect data, you must install VRAdvisor on the system for which you are collecting data. If you do not plan to install VRAdvisor on the system, you can use the data collection script to collect data and use VRAdvisor on another system to analyze the results.

On UNIX, collect the sample of data using one of the following methods:

- [Collecting data using the VRAdvisor wizard](#)
- [Collecting data using the vxstat command](#)
- [Collecting data using the data collection script](#)

Prerequisite

- If you use Veritas Volume Manager (VxVM) volumes, make sure that you import the disk group containing the required volumes onto your system.

Supported locales

VRAdvisor requires the data to be collected in a supported locale. Before using any of the methods for collecting data, set the locale to a supported locale. For UNIX, VRAdvisor supports the following locales:

English:

C
en_US.UTF-8
ios_8859_1

Japanese:

ja
ja_JP.PCK
ja_JP.UTF-8
ja_JP.eucJP

Any of the methods of data collection include a date stamp with each sample of data. When you use VRAdvisor to analyze the sample data file, VRAdvisor uses the date stamp to determine the data collection interval. To enable VRAdvisor to automatically determine the data collection interval, the date must be displayed in one of the following formats.

Date formats for English locales:

Fri Oct 1 14:37:13 2004
Fri Oct 1 14:37:13 PDT 2004
Friday October 1 17:37:13 PDT 2004

Date formats for Japanese locales:

2004年09月24日16時21分23秒

2004年09月24日(金)16時20分59秒

If VRAdvisor cannot determine the data collection interval, it prompts you to specify the data collection interval.

Collecting data using the VRAdvisor wizard

The VRAdvisor wizard supports collecting data for all disks or disk groups on the host on which VRAdvisor is installed.

To collect data using the VRAdvisor wizard

- 1 Set the locale to a supported locale. For example:
export LC_ALL=C
- 2 Change directory as follows:
cd /opt/VRTSvradv/bin
- 3 Launch the VRAdvisor wizard on Solaris, using the following command:
./vradvgui
- 4 In the Welcome page, select **Data Collection**, and then click **Next**.
- 5 Complete the Data Collection page as follows:

File Name	<p>Enter the name of the file where the data write samples will be collected.</p> <p>Make sure the name is not being used by another application.</p> <p>If a file already exists with that filename, or if the path is incorrect, a message is displayed.</p>
Duration for which data is to be collected	<p>Enter the duration in days or hours. The default value is 14 days. The maximum duration is 30 days.</p>
Interval	<p>Enter a value, in seconds, to indicate the frequency at which you want the data to be collected. The default value is 120 seconds.</p>
Details	<p>If you have VxVM installed, select the appropriate disk group in Disk Group. If you do not have VxVM installed, the Disk Group selection is not available.</p> <p>Select the required volumes individually, or click Select All to select all of the available volumes in the selected disk group.</p>

- 6 Click **Next**. The Confirmation message appears.
- 7 To start the data collection process immediately, click **Yes**. To go back and make any changes, click **No**.
- 8 The Data Collection Summary page indicates that the data collection has started. It also displays a summary of the specifications you entered for the data collection.
- 9 Click **Finish**. VRAdvisor continues to collect data for the specified duration, although the wizard window closes.
The data collection wizard displays an appropriate error message if it is unsuccessful in starting the data collection process. Select **Cancel**, fix the reported error and launch the data collection wizard again.

After the data collection completes, the file specified by **File Name** contains the sample of data in a format that can be used for analysis by VRAdvisor. To proceed, see “[Analyzing the sample of data](#)” on page 27.

Collecting data using the vxstat command

If you do not want to install VRAdvisor and VxVM is installed on your system, use the `vxstat` command to collect data.

To collect data using vxstat

- 1 Set the locale to a supported locale. For example:

```
# export LC_ALL=C
```
- 2 To collect the data in the format required for analysis, use the following command with exactly the parameters shown:

```
# vxstat -g dgname -i interval -c count volumes > filename
```

where:

interval is the data collection interval in seconds.

count is the number of samples to collect.

volumes is a list of volume names separated by spaces.

For example, use the following command to collect a sample of data every 120 seconds and to collect 5040 samples. The volumes are the data volumes `hr_dv01` and `hr_dv02` in the disk group `hrdg`. The results are collected in the file `vra_in`.

```
# vxstat -g hrdg -i 120 -c 5040 hr_dv01 hr_dv02 > vra_in
```

After the data collection completes, the file *filename* (for example, *vra_in*) contains the sample of data in the vxstat format, which can be used for analysis by VRAdvisor. To proceed, see “[Analyzing the sample of data](#)” on page 27.

Collecting data using the data collection script

To collect sample data for analysis by VRAdvisor, you need not have VRAdvisor installed on the system. Use the data collection script to collect data for AIX, HP-UX, Linux, or Solaris systems.

The data collection script:

- collects sample data using the appropriate method for the host. If VxVM is installed, the script uses the `vxstat` command to collect data. If VxVM is not installed, the script uses the appropriate operating system command to collect data:
 - AIX: uses `lvmstat` to collect data for all logical volumes.
 - HP-UX, Linux: uses `sar` to collect data for all disks.
 - Solaris: uses `iostat` to collect data for all disks.
- stores the collected data and its associated metadata in a format that can be analyzed by VRAdvisor.
- sends notifications by email when an error occurs or when the data collection has successfully completed.

To collect data using the script

- 1 Set the locale to a supported locale. For example:

```
# export LC_ALL=C
```
- 2 Copy the data collection script and all required files to the system on which you plan to collect data, from one of the following locations:
 - `volume_replicator/tools/vradvisor/scripts/` directory of the Veritas product disc
 - `scripts` directory of a Solaris system where VRAdvisor is installed. By default, the `/opt/VRTSvradv/scripts` directory.
 - `scripts` folder of a Windows system where VRAdvisor is installed. By default, the folder `Program Files/VERITAS/Volume Replicator Advisor/scripts`
 - **Note:** Make sure to copy all the files in the directory.
- 3 To collect the data in the format required for analysis, use the following script:

```
# sh vra_data_collection.sh [-g dgname] [-i interval] \  
[-t duration] [-d dirname] [-v volume_list] \  
[-m maillist]
```

where:

dgname is a comma-separated list of disk groups for which data needs to be collected, if VxVM is installed. The default is all disk groups on the host.

interval is the data collection interval in seconds. The default is 120 seconds.

duration is the duration for which to collect data, in hours or days. Specify the suffix "h" for hours, "d" for days. The default is 14 days.

volume_list is a comma-separated list of volume names for which data needs to be collected, if VxVM is installed. The default is all volumes on the host.

dirname is the directory in which to store the collected data files. If no directory is specified, the files are stored in the current working directory.

maillist is a comma-separated list of email-ids to receive notifications.

By default, the script sends an email notification to the root user of the host.

After the data collection completes, the collected data file contains the sample of data in a format that can be used for analysis by VRAdvisor. The file names are as follows:

- *hostname_dgname_timestamp.vxstat*
Indicates the file contains *vxstat* output. This is the default if VxVM is installed on the system where data is collected. Any associated metadata is in a file with the same name plus the extension *.meta*.
- *hostname_dgname_timestamp.vra*
Indicates that the data is in VRAdvisor CSV format. The script converts the data to this format if a command other than *vxstat* was used to collect the data. Any associated metadata is in a file with the same name plus the extension *.meta*.

To proceed, see “[Analyzing the sample of data](#)” on page 27.

To show help for the script options

```
◆ # sh vra_data_collection.sh -h
```

Examples

Example 1

Collect data every minute for 30 hours.

```
# sh vra_data_collection.sh -i 60 -t 30h
```

Example 2

Store the collected data in directory *samples* and notify *abc@example.com*.

```
# sh vra_data_collection.sh -d samples -m abc@example.com
```

Example 3

Collect data for the VxVM disk groups *dg1* and *dg2*.

```
# sh vra_data_collection.sh -g dg1,dg2
```

Example 4

Collect data for the VxVM volumes `vol1`, `vol2`, and `vol3` in the disk group `dg1`.

```
# sh vra_data_collection.sh -g dg1 -v vol1,vol2,vol3
```

Example 5

Collect data for the VxVM volumes `vol1` and `vol2` in the disk group `dg1`, every 5 minutes, for 15 days. Store collected data in directory `samples`. Notify `abc@example.com` and `pqr@example.com`.

```
# sh vra_data_collection.sh -d samples -i 300 -d 15d -g dg1 \  
-v vol1,vol2 -m abc@example.com,pqr@example.com
```

Collecting the sample of data on Windows

VRAdvisor can be used to collect and analyze a sample of data. You can collect data using the VRAdvisor wizard or the `diskStats` command. To use VRAdvisor to collect data, you must install VRAdvisor on your system. If you do not plan to install VRAdvisor on your system, use the `diskStats` command to collect data.

On Windows, collect the sample of data using one of the following methods:

- [Collecting the sample of data using the VRAdvisor wizard](#)
- [Collecting the sample of data using the diskStats command](#)

Prerequisite

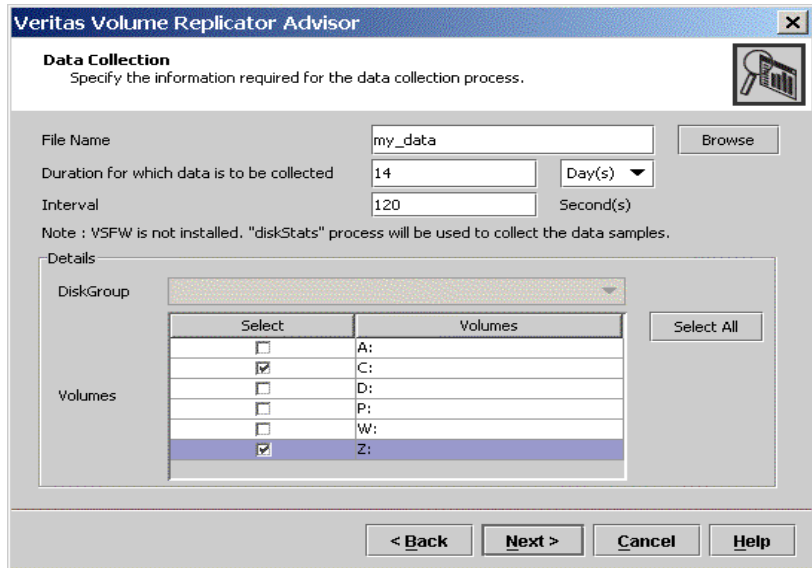
- If you are using VSFV volumes, then ensure that you import the disk group containing the required volumes onto your system.

Collecting the sample of data using the VRAdvisor wizard

To collect data using the VRAdvisor wizard

- 1 To launch the VRAdvisor wizard on Windows, select **Start > Programs > VERITAS > Volume Replicator Advisor > VRAdvisor Wizard**.
- 2 In the Welcome page, select **Data Collection** and then click **Next**.
The Data Collection page appears.

Note: On Windows, only the `diskStats` command will be used to collect data.



3 Complete the Data Collection page as follows:

- | | |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| File Name | Enter the name of the file where the data write samples will be collected.

Make sure the name is not being used by another application.

If a file already exists with that filename, or if the path is incorrect, a message is displayed. |
| Duration for which data is to be collected | Enter the duration in days or hours. The default value is 14 days. The maximum duration is 30 days. |
| Interval | Enter a value, in seconds, to indicate the frequency at which you want the data to be collected. The default value is 120 seconds. |

Details

Select the required volumes individually, or click **Select All** to select all of the available volumes in the selected disk group.

Only volumes with drive letters are displayed.

On Windows, the Disk Group field is not available.

- 4 Click **Next**. The Confirmation message appears.
- 5 To start the data collection process immediately, click **Yes**. To go back and make any changes, click **No**.
The Data Collection Summary page indicates that the data collection has started. It also displays a summary of the specifications you entered for the data collection.
- 6 Click **Finish**. VRAadvisor continues to collect data for the specified duration, although the wizard window closes.
The data collection wizard displays an appropriate error message if it is unsuccessful in starting the data collection process. Select **Cancel**, fix the reported error and launch the data collection wizard again.

After the data collection completes, the file specified by **File Name** contains the sample of data in a format that can be used for analysis by VRAadvisor. To proceed, see “[Analyzing the sample of data](#)” on page 27.

Collecting the sample of data using the diskStats command

On Windows, use the `diskStats` command to collect the data required for analysis. This command can be used to collect data whether or not the Veritas Storage Foundation is installed on the system. The `diskStats` utility is installed in the following location:

```
\VERITAS\Volume Replicator Advisor\bin\diskStats.exe
```

To collect data using the diskStats command

- 1 Navigate to the specified path:

```
\VERITAS\Volume Replicator Advisor\bin\
```

- 2 At the prompt, enter the following command with exactly the parameters shown:

Note: The `diskStats` command can accept only drive letters of the volumes as inputs. Volume names or mount points are not supported. Volumes created by any application are supported.

```
diskStats [-i interval [-c count]] \  
          <drive 1> [[drive 2][drive 3]... ]
```

The command will display the output on the console.

To save the output to a file, you can redirect the output to a named file using the command:

```
diskStats [-i interval [-c count]] \  
          <drive 1> [[drive 2][drive 3]... ] > <filename>
```

After the data collection completes, the file `filename` contains the sample of data in the `diskStats` format, which can be used for analysis by VRAdvisor. To proceed, see [“Analyzing the sample of data”](#) on page 27.

Analyzing the sample of data

You can use VRAdvisor to analyze the sample of data that you have collected. VRAdvisor analyzes the sample of data according to parameters that you specify such as available network bandwidth and network outage. In addition, VRAdvisor enables you to perform a What-If analysis by varying the values of the parameters. The output of the analysis gives the network bandwidth required to replicate in synchronous mode and the SRL size required for a given bandwidth and for the given outages to replicate in asynchronous mode. The results of the analysis help you to set up an optimum configuration for VVR. For more information on some of the considerations and formulas used in determining the size of the SRL, see [Appendix A, “Sizing the SRL” on page 39](#). VRAdvisor enables you to analyze data collected on any of the supported platforms described in [“Collecting the sample of data” on page 15](#). However, to analyze the data, you must install and use VRAdvisor on either a Solaris or a Windows system.

Prerequisites

- All the files to be analyzed must be present in a single directory.
- The sample of data must be available in a format required by VRAdvisor. VRAdvisor accepts the following formats:
 - `vxstat` output
 - `diskStats` output
 - VRAdv CSV format (used by VRAdvisor wizard or the UNIX data collection script)

Analyzing the collected data

To analyze the collected data using VRAdvisor

- 1 Set the locale to a supported locale. For example:

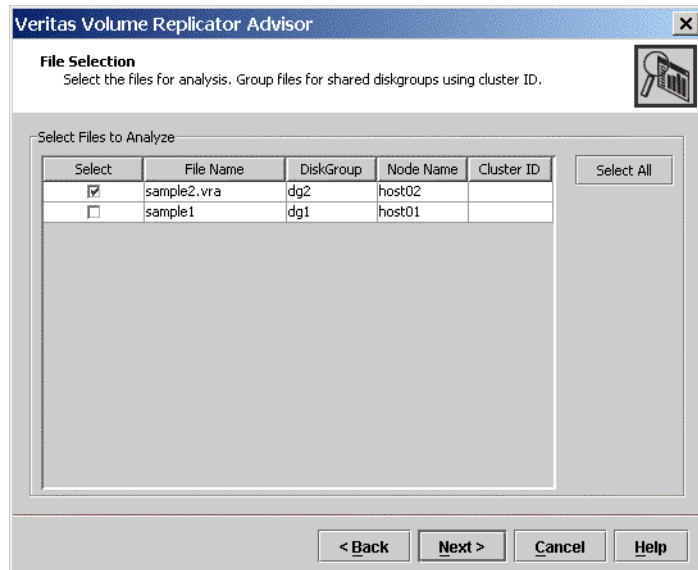
```
# export LC_ALL=C
```
- 2 Launch the VRAdvisor wizard.
On Solaris:
 - a Change directory as follows:

```
# cd /opt/VRTSvradv/bin
```
 - b Launch the VRAdvisor wizard, using the following command:

```
# ./vradvgui
```On Windows:  
Choose **Start > Programs > VERITAS > Volume Replicator Advisor > VRAdvisor Wizard**.
- 3 In the Welcome page, select **Analysis**, and then click **Next**.
- 4 In the Directory Specification page, enter the name of the directory containing the data files to be analyzed. All of the files to be analyzed must be in the same directory.
The specified directory must contain the data files and any metadata files associated with each data file. The associated metadata and data files must have the same name except for the extension. Metadata files must have the extension `.meta`.


- 5 In the File Selection page, VRAdvisor displays the list of files in a table. Select the files to be analyzed.

Note: Files containing information from nodes that will use the same network bandwidth for replication should be analyzed together. Otherwise, files should not be selected together. In order for the files to be analyzed together, the data collection for each node must start at the same time.



- To analyze files for nodes within a shared disk group, specify a cluster ID for each shared disk group. The analysis wizard uses the cluster ID field to group the files. During the analysis, I/O statistics for files sharing the same cluster ID are considered together.
- Enter the DiskGroup name, Node Name and Cluster ID, if necessary.

- 6 In the Block Size and Collection Interval Specification page, specify the metadata as follows:



If the data was collected using the data collection script for UNIX platforms, the generated files contain metadata such as block size, and data collection interval.

If the files do not contain metadata, because the data was collected using operating system commands or the VRAdvisor Wizard, enter the appropriate metadata:

- Specify the block size, if required.
 - If no timestamps are present in the file, or if VRAdvisor is unable to parse the timestamps, specify the interval used during the data collection.
- 7 In the Volume or Disk Selection page, select the tab for each selected file. For each file, the wizard lists the disks or volumes for which data has been collected.
When selecting disks or volumes, make sure you do not select:
 - RAID-5 volumes because these are not supported.
 - Sub-level volumes (if the volumes are layered volumes). Select only the top-level volumes.
 - The volume that you intend to use as the SRL.

- Drives or volumes containing high-activity data that is not be replicated.
Using VRA to analyze data from drives or volumes containing high activity data that is not be replicated will lead to erroneous results. Select the volumes or disks to be analyzed, and then Click **Next**.
- 8 The RVG Summary page displays the disks or volumes that were selected for analysis. The disks or volumes for each analyzed file are grouped under an RVG name.
Click **Back** to modify the selections, or click **Next** to continue.
- 9 In the Network Parameters for Analysis page, specify the parameters that apply to all defined RVGs.
 - **Network Bandwidth Available for Replication** indicates the total bandwidth of the network across which you are replicating. Enter the network bandwidth that will be available for replication. Select the unit for the network bandwidth from the drop-down list. The default is 100 Mbps.

Note: Before specifying the network bandwidth you must also consider the loss of available bandwidth because of the TCP-IP/UDP headers, because VRAdvisor does not handle this.

- **Network Outage Duration** indicates the maximum expected outage times applicable for all defined RVGs, for example, the time during which the network link is unavailable for the network that is used by all of the RVGs for replication. Enter the duration of the network outage in days, hours, or minutes. The default is zero.
Click **Next**.
- 10 Complete the RVG Specific Parameters page, as follows:
For each RVG, select the tab and specify the following parameters:
 - **Bandwidth Limit** indicates the bandwidth throttling for that RVG. The default is 0 (zero), which indicates that no bandwidth limit applies.
 - **Secondary Outage Duration** indicates the maximum expected outage times specific to that RVG, for example, the time during which the Secondary host for the RVG is unavailable. Enter the outage duration in days, hours, or minutes. The default is one hour.
 - **Apply to all RVG(s)** indicates that the same bandwidth limit and outage duration apply to all RVGs. Select this check box to enable the All tab and disable the RVG-specific tabs.
Click **Next**.

The Summary of Inputs page displays. The Total Outage Duration column shows the sum of the Network Outage Duration and the Secondary Outage for that RVG.

- 11 In the Summary of Inputs page, click **Back** to modify the parameters, or select **Analyze** to start the analysis.
VRAdvisor displays the results of the analysis for the selected data files.

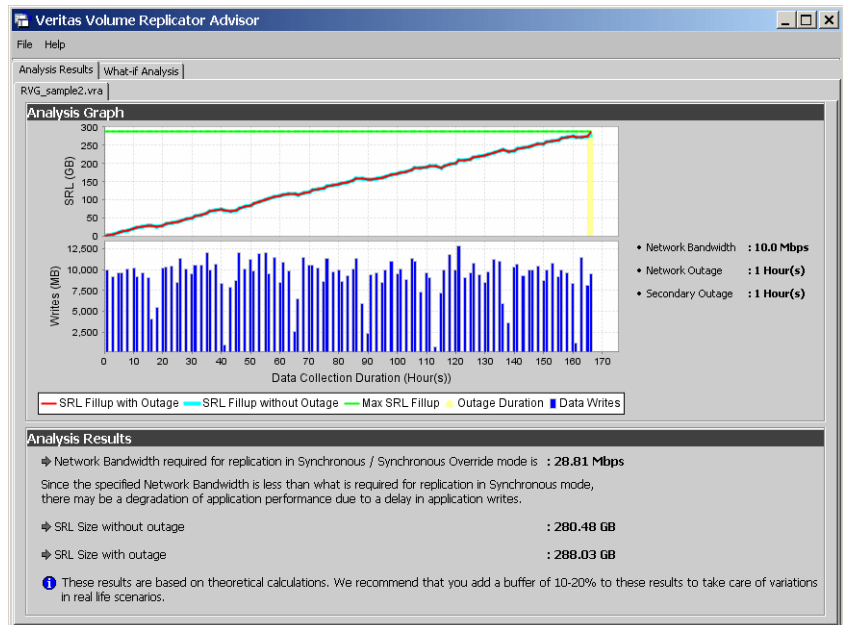
Understanding the results of the analysis

After the analysis completes, VRAdvisor displays the results of the analysis and also enables you to recalculate the results after changing some parameters. The results are described in the following sections:

- [Viewing the analysis results](#)
- [Recalculating the analysis results](#)
- [Recording and viewing the results](#)

Viewing the analysis results

After the analysis completes, the Analysis Results page displays by default.



The Analysis Results page displays the results of the analysis for each RVG. Select the tab for the RVG to display the results for that RVG. The results pages display the following information:

Analysis graph

The Analysis Graph section shows the following information:

- The top graph shows the SRL fillup in megabytes (MB) on the y-axis. The fillup rate is shown both with the specified outages and with no outage. The

x-axis shows the data write duration values. The peak SRL fillup size is shown against a max outage window, displayed in yellow, which indicates the worst case scenario.

Note: If the SRL fillup value in the graph is steadily increasing, with the maximum at the last data write duration, it indicates that you do not have sufficient network bandwidth for the amount of data writes in the sample of data.

- The bar graph shows the values of the Application Writes in bytes for the y-axis. The x-axis shows the data write duration values.
- To the right of the graphs, the page displays the values you specified for network bandwidth and the outage parameters.

Analysis results

The Analysis Results section displays the following information:

- Network bandwidth required for synchronous replication.
If the required bandwidth is more than the bandwidth that you specified, then VRAdvisor displays a message to indicate that the performance of the application writing to the disk writes will be affected.
- The required SRL size with the specified outage.
- The required SRL size with no outage.

Note: We recommend that you add a 10-20 percent buffer to the values calculated by VRAdvisor when setting up your VVR configuration. VRAdvisor analyzes the data based on the specified values, which could be affected by factors that VRAdvisor does not consider, such as TCP/IP headers overhead, or network congestion.

Recalculating the analysis results

You can recalculate the analysis results in the following ways:

- [Applying different parameters to the existing sample of data](#)
- [Performing What-if analysis](#)

Applying different parameters to the existing sample of data

You can recalculate the analysis results by changing the values you specified for the network bandwidth and the outage durations.

To recalculate the analysis results

- 1 To change the values you specified, select **File > Change Inputs**.
- 2 On the Network Parameters for Analysis page, specify new values for any of the fields as required. Click **Next** to specify RVG-specific parameters, or click **Back** to change volume or disk selection.
- 3 Continue using the **Next** and **Back** buttons to navigate through the input pages and change values as required. See [“Analyzing the collected data”](#) on page 28 for information.
- 4 When the values are correct, click **Next** to navigate to the Summary of Inputs page.
- 5 Select **Analyze** to start the analysis.
VRAdvisor performs the analysis of the data using the changed values and displays the results.

Performing What-if analysis

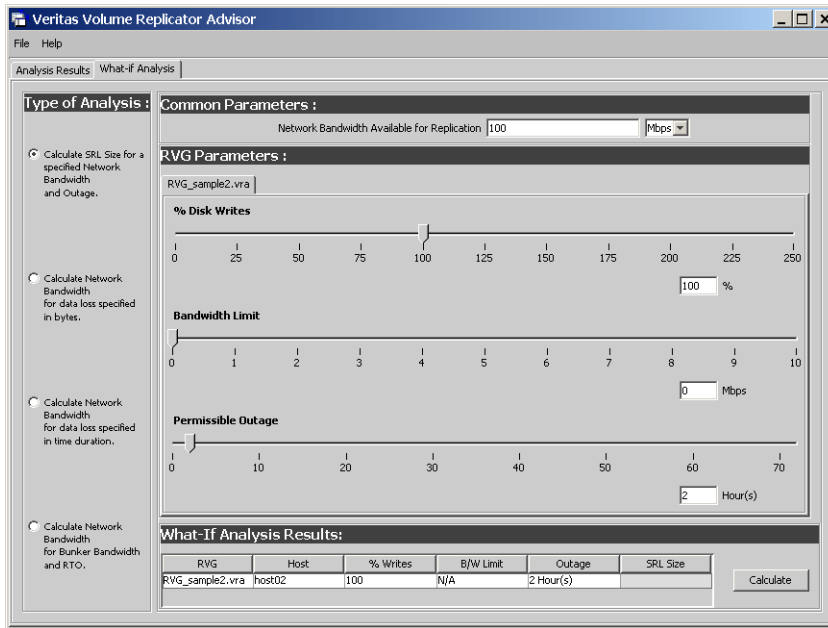
After checking the analysis results, you can use the What-if Analysis page to do additional calculations, to plan for future requirements or alternative scenarios. You can vary the parameters and recalculate the results according to different criteria. For example, you can vary the network bandwidth parameter to see what effect it would have on the SRL size. Or, you can specify a potential SRL size and see how much network bandwidth would be required for that SRL size.

Note: Before specifying the network bandwidth, you must also consider the loss of available bandwidth due to the TCP-IP/UDP headers, because VRAdvisor cannot manage this.

What-if Analysis also enables you to vary the percentage of disk writes as compared to the sample of data that was analyzed. For example, if you anticipate that your future needs will involve twenty percent more disk writes, set the percentage of disk writes to 120% and recalculate.

To recalculate results using the What-If Analysis

- 1 Select the What-If Analysis tab.



- 2 Select the appropriate option on the left side of the What-If Analysis page to recalculate results, as follows:

- **Calculate SRL Size for a specified Network Bandwidth and Outage.**
Use this option to calculate the SRL size for a specified network bandwidth and outage duration.
Available parameters for this option are % Disk Writes and Permissible Outage.
- **Calculate the Network Bandwidth for data loss specified in bytes.**
Use this option to calculate the network bandwidth that would be required to minimize the amount of data loss at the Primary host.
Available parameters for this option are % Disk Writes and Data loss in bytes.
- **Calculate Network Bandwidth for data loss specified in time duration.**
Use this option to calculate the network bandwidth that would be required to minimize the amount of data loss at the Primary host.
Available parameters for this option are % Disk Writes and Data loss in time.
- **Calculate Network Bandwidth for Bunker and RTO.**

In a Bunker replication setup, the available bandwidth determines the RPO and the RTO that can be achieved after a disaster. Use this option to calculate the bandwidth required between the Primary and the Secondary, and between the Bunker and the Secondary, based on the desired RPO and RTO.

Available parameters for this option are % Disk Writes and RTO. The Have Bunker check box indicates whether the RVG has a bunker attached.


The right side of the page displays the parameters you can specify for each option, and the corresponding slider bars.

- 3 In the Common Parameters section, change the bandwidth value shared by all RVGs.
- 4 In the RVG Parameters section, select the tab for the RVG that you want to change, and then use the slider bar to specify the value for each parameter. Each slider has a default range of values, which can be customized using the **Preferences** option that is available from the **File** menu. For more information, see “[Changing the value ranges on the slider bar](#)” on page 37.
- 5 Click **Calculate** at the lower region of the page. The What-if Analysis Results are displayed in this section.

Changing the value ranges on the slider bar

Follow the steps given below to change the value ranges for the slider bars.

- 1 Make sure the option for which you want to change the value ranges is selected in the left side of the What-if Analysis page.
- 2 Use the **File > Preferences** option to display the Preferences page.



Note: The Preferences dialog box displays parameters corresponding to the calculate option that you selected.

- 3 Change the values on the Preferences page as required:
 - Select the Unit for each option from the drop-down box.
 - Specify the appropriate values in the **Maximum** and **Minimum** fields. These values are used to indicate the range of values available on the slider bar.
- 4 Click **Ok**.

Recording and viewing the results

VRAdvisor records any values that you had specified in the analysis phase and the results of the What-if Analysis to a file, which uses the following naming convention:

`VRAdvResults_Datestamp_and_Timestamp.txt`

For Windows, the file is located in the `VERITAS/Volume Replicator Advisor/results` subdirectory. For Solaris, the file is located in the `/opt/VRTSvradv/results` subdirectory.

Every time you start the Analysis wizard, this file is automatically created and can be referenced later.

Sizing the SRL

The size of the SRL is critical to the performance of replication. You can use VRAdvisor to help determine the appropriate SRL size. This section describes some of the considerations in determining the size of the SRL. VRAdvisor uses the formulas described in this section to determine the appropriate SRL size.

Overview

When the SRL overflows for a particular Secondary, the RLINK corresponding to that Secondary is marked `STALE` and 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 avoid SRL overflows. The SRL size needs to be large enough to satisfy four constraints:

- It must not overflow for asynchronous RLINKs during periods of peak usage when replication over the RLINK may fall far behind the application.
- It must not overflow while a Secondary RVG is being synchronized.
- It must not overflow while a Secondary RVG is being restored.
- It must not overflow during extended outages (network or Secondary node).

Note: The size of the SRL must be at least 110 MB. If the size that you have specified for the SRL is less than 110 MB, VVR displays an error message which prompts you to specify a value that is equal to or greater than 110 MB.

To determine the size of the SRL, you must determine the size required to satisfy each of these constraints individually. Then, choose a value at least equal to the maximum so that all constraints are satisfied. The information needed to perform this analysis, presented below, includes:

- The maximum expected downtime for Secondary nodes

- The maximum expected downtime for the network connection
- The method for synchronizing Secondary data volumes with data from Primary data volumes. If the application is shut down to perform the synchronization, the SRL is not used and the method is not important. Otherwise, this information could include: the time required to copy the data over a network, or the time required to copy it to a tape or disk, to send the copy to the Secondary site, and to load the data onto the Secondary data volumes.

Note: If the Automatic Synchronization option is used to synchronize the Secondary, the previous paragraph is not a concern.

If you are going to perform Secondary backup to avoid complete resynchronization in case of Secondary data volume failure, the information needed also includes:

- The frequency of Secondary backups
- The maximum expected delay to detect and repair a failed Secondary data volume
- The expected time to reload backups onto the repaired Secondary data volume

Peak usage constraint

For some configurations, it might be common for replication to fall behind the application during some periods and catch up during others. For example, an RLINK might fall behind during business hours and catch up overnight if its peak bandwidth requirements exceed the network bandwidth. Of course, for synchronous RLINKs, this does not apply, as a shortfall in network capacity would cause each application write to be delayed, so the application would run more slowly, but would not get ahead of replication.

For asynchronous RLINKs, the only limit to how far replication can fall behind is the size of the SRL. If it is known that the peak write rate requirements of the application exceed the available network bandwidth, then it becomes important to consider this factor when sizing the SRL.

Assuming that data is available providing the typical application write rate over a series of intervals of equal length, it is simple to calculate the SRL size needed to support this usage pattern:

- 1 Calculate the network capacity over the given interval (BW_N).

- For each interval n , calculate SRL log volume usage (LU_n), as the excess of application write rate (BW_{AP}) over network bandwidth ($LU_n = BW_{AP(n)} - BW_N$).

Note: In a shared environment, you must consider the write rates on all the nodes in the cluster. The application write rate (BW_{AP}) should reflect the aggregate of the write rates on each node.

- For each interval, accumulate all the SRL usage values to find the cumulative SRL log size (LS):

$$LS_n = \sum_{i=1..n} LU_i$$

The largest value obtained for any LS_n is the value that should be used for SRL size as determined by the peak usage constraint. See the table “[Example calculation of SRL size required to support peak usage period](#)” on page 41, which shows an example of this calculation. The third column, Application, contains the maximum likely application write rate per hour. The fourth column, Network, shows the network bandwidth. The fifth column, SRL Usage, shows the difference between application write rate and network bandwidth obtained for each interval. The sixth column, Cumulative SRL Size, shows the cumulative difference every hour. The largest value in column 6 is 37 gigabytes. The SRL should be at least this large for this application.

Note that several factors can reduce the maximum size to which the SRL can fill up during the peak usage period. Among these are:

- The `latencyprot` characteristic can be enabled to restrict the amount by which the RLINK can fall behind, slowing down the write rate.
- The network bandwidth can be increased to handle the full application write rate. In this example, the bandwidth should be 15 gigabytes/hour—the maximum value in column three.

Table A-1 Example calculation of SRL size required to support peak usage period

| Hour Starting | Hour Ending | Application (GB/hour) | Network (GB/hour) | SRL Usage (GB) | Cumulative SRL Size (GB) |
|---------------|-------------|-----------------------|-------------------|----------------|--------------------------|
| 7am | 8 a.m. | 6 | 5 | 1 | 1 |
| 8 | 9 | 10 | 5 | 5 | 6 |

Table A-1 Example calculation of SRL size required to support peak usage period

| Hour Starting | Hour Ending | Application (GB/hour) | Network (GB/hour) | SRL Usage (GB) | Cumulative SRL Size (GB) |
|---------------|-------------|-----------------------|-------------------|----------------|--------------------------|
| 9 | 10 | 15 | 5 | 10 | 16 |
| 10 | 11 | 15 | 5 | 10 | 26 |
| 11 | 12 p.m. | 10 | 5 | 5 | 31 |
| 12 p.m. | 1 | 2 | 5 | -3 | 28 |
| 1 | 2 | 6 | 5 | 1 | 29 |
| 2 | 3 | 8 | 5 | 3 | 32 |
| 3 | 4 | 8 | 5 | 3 | 35 |
| 4 | 5 | 7 | 5 | 2 | 37 |
| 5 | 6 | 3 | 5 | -2 | 35 |

Note: In a shared environment, the values in the Application column should include write rates on all the nodes. For example, if in one hour, the write rate on `seattle1` is 4 GB and the write rate on `seattle2` is 2 GB, the application write rate is 6 GB/hour.

Synchronization period constraint

When a new Secondary is added to an RDS, its data volumes must be synchronized with those of the Primary unless the Primary and the Secondary data volumes have been zero initialized and the application has not yet been started. You also need to synchronize the Secondary after a Secondary data volume failure, in case of SRL overflow, or after replication is stopped.

This section applies if you choose *not* to use the automatic synchronization method to synchronize the Secondary. Also, this constraint does not apply if you choose to use a method other than automatic synchronization and if the application on the Primary can be shut down while the data is copied to the Secondary. However, in most cases, it might be necessary to synchronize the Secondary data volumes with the Primary data volumes while the application is still running on the Primary. This is performed using one of the methods described in the *Veritas Volume Replicator Administrator's Guide*.

During the synchronization period, the application is running and data is accumulating in the SRL. If the SRL overflows during the process of synchronization, the synchronization process must be restarted. Thus, to ensure that the SRL does not overflow during this period, it is necessary that the SRL be sized to hold as much data as the application writes during the synchronization period. After starting replication, this data is replicated and the Secondary eventually catches up with the Primary.

Depending on your needs, it may or may not be possible to schedule the synchronization during periods of low application write activity. If it is possible to complete the synchronization process during a period of low application write activity, then you must ensure that the SRL is sized such that it can hold all the incoming writes during this period. Otherwise, the SRL may overflow. Using VRAdvisor enables you to arrive at an optimum SRL size.

Secondary backup constraint

VVR provides a mechanism to perform periodic backups of the Secondary data volumes. In case of a problem that would otherwise require a complete resynchronization using one of the methods described in “[Synchronization period constraint](#)” on page 42, a Secondary backup, if available, can be used to bring the Secondary online much more quickly.

A Secondary backup is made by defining a Secondary checkpoint and then making a raw copy of all the Secondary data volumes. Should a failure occur, the Secondary data volumes are restored from this local copy, and then replication proceeds from the checkpoint, thus replaying all the data from the checkpoint to the present.

The constraint introduced by this process is that the Primary SRL must be large enough to hold all the data logged in the Primary SRL after the creation of the checkpoint corresponding to the most recent backup. This depends largely on three factors:

- The application write rate.
- The frequency of Secondary backups.

Thus, given an application write rate and frequency of Secondary backups, it is possible to come up with a minimal SRL size. Realistically, an extra margin should be added to an estimate arrived at using these figures to cover other possible delays, including:

- Maximum delay before a data volume failure is detected by a system administrator.
- Maximum delay to repair or replace the failed drive.
- Delay to reload disk with data from the backup tape.

To arrive at an estimate of the SRL size needed to support this constraint, first determine the total time period the SRL needs to support by adding the period planned between Secondary backups to the time expected for the three factors mentioned above. Then, use the application write rate data to determine, for the worst case, the amount of data the application could generate over this time period.

Note: Even if only one volume failed, all volumes must be restored.

Secondary downtime constraint

When the network connection to a Secondary node, or the Secondary node itself, goes down, the RLINK on the Primary node detects the broken connection and responds. If the RLINK has its `synchronous` attribute set to `fail`, the response is to fail all subsequent write requests until the connection is restored. In this case, the SRL does not grow, so the downtime constraint is irrelevant. For all other types of RLINKs, incoming write requests accumulate in the SRL until the connection is restored. Thus, the SRL must be large enough to hold the maximum output that the application could be expected to generate over the maximum possible downtime.

Maximum downtimes may be difficult to estimate. In some cases, the vendor may guarantee that failed hardware or network connections will be repaired within some period. Of course, if the repair is not completed within the guaranteed period, the SRL overflows despite any guarantee, so it is a good idea to add a safety margin to any such estimate.

To arrive at an estimate of the SRL size needed to support this constraint, first obtain estimates for the maximum downtimes which the Secondary node and network connections could reasonably be expected to incur. Then, use the application write rate data to determine, for the worst case, the amount of data the application could generate over this time period. With the introduction of the `autodcm` mode of SRL overflow protection, sizing the SRL for downtime is not essential to prevent SRL overflow because the changed blocks are no longer stored in the SRL. However, note that the Secondary is inconsistent during the replay of the DCM, and hence it is still important for the SRL to be large enough to cover most eventualities.

Additional factors

Once estimates of required SRL size have been obtained under each of the constraints described above, several additional factors must be considered.

For the synchronization period, downtime and Secondary backup constraints, it is not unlikely that any of these situations could be immediately followed by a

period of peak usage. In this case, the Secondary could continue to fall further behind rather than catching up during the peak usage period. As a result, it might be necessary to add the size obtained from the peak usage constraint to the maximum size obtained using the other constraints. Note that this applies even for synchronous RLINKs, which are not normally affected by the peak usage constraint, because after a disconnect, they act as asynchronous RLINKs until caught up.

Of course, it is also possible that other situations could occur requiring additions to constraints. For example, a synchronization period could be immediately followed by a long network failure, or a network failure could be followed by a Secondary node failure. Whether and to what degree to plan for unlikely occurrences requires weighing the cost of additional storage against the cost of additional downtime caused by SRL overflow.

Once an estimate has been computed, one more adjustment must be made to account for the fact that all data written to the SRL also includes some header information. This adjustment must take into account the typical size of write requests. Each request uses at least one additional disk block for header information.

For AIX, Linux, and Solaris, the adjustments are as follows:

| If Average Write Size is: | Add This Percentage to SRL Size: |
|----------------------------------|-----------------------------------------|
| 512 bytes | 100% |
| 1K | 50% |
| 2K | 25% |
| 4K | 15% |
| 8K | 7% |
| 10K | 5% |
| 16K | 4% |
| 32K or more | 2% |

For HP-UX, the adjustments are as follows:

| If Average Write Size is: | Add This Percentage to SRL Size: |
|----------------------------------|-----------------------------------------|
| 1K | 100% |
| 2K | 50% |
| 4K | 25% |

| If Average Write Size is: | Add This Percentage to SRL Size: |
|----------------------------------|-----------------------------------------|
| 8K | 13% |
| 10K | 10% |
| 16K | 6% |
| 32K or more | 3% |

Example

This section shows how to calculate the SRL size for a VVR configuration. First, collect the relevant parameters for the site. For this site, they are as follows:

| | |
|-------------------------------------|----------------------|
| Application peak write rate | 1 gigabyte/hour |
| Duration of peak | 8 am - 8 pm |
| Application off-peak write rate | 250 megabytes/hour |
| Average write size | 2 kilobytes |
| Number of Secondary sites | 1 |
| Type of RLINK | synchronous=override |
| Synchronization Period: | |
| application shutdown | no |
| copy data to tape | 3 hours |
| send tapes to Secondary site | 4 hours |
| load data | 3 hours |
| Total | 10 hours |
| Maximum downtime for Secondary node | 4 hours |
| Maximum downtime for network | 24 hours |
| Secondary backup | not used |

Because synchronous RLINKs are to be used, the network bandwidth must be sized to handle the peak application write rate to prevent the write latency from growing. Thus, the peak usage constraint is not an issue, and the largest constraint is that the network could be out for 24 hours. The amount of data accumulating in the SRL over this period would be:

(Application peak write rate x Duration of peak) +
(Application off-peak write rate x Duration of off peak).

In this case, the calculation would appear as follows:

$$1 \text{ GB/hour} \times 12 \text{ hours} + 1/4 \text{ GB/hour} \times 12 = 15 \text{ GB}$$

An adjustment of 25% is made to handle header information. Since the 24-hour downtime is already an extreme case, no additional adjustments are needed to handle other constraints. The result shows that the SRL should be at least 18.75 gigabytes.

Index

A

- Add or Remove Programs option 13
- analysis results
 - recalculating 34
 - saving 38
 - viewing from file 38
- analyzing data
 - overview 27
 - performing what-if analysis 35
- application writes, displaying 34

B

- backup constraints, Secondary 43
- backup interval, see Secondary outage 36
- bandwidth, required for synchronous replication 34

C

- collecting data
 - on UNIX 16
 - on Windows 22
- constraints
 - peak usage 40
 - Secondary backup 43
 - Secondary downtime 44
 - synchronization period 42

D

- data analysis results
 - recalculating 34
 - saving 38
 - viewing from file 38
 - viewing in VRAdvisor 33
- data analyzing
 - overview 27
 - performing what-if analysis 35
- data collection
 - on UNIX 16
 - on Windows 22

- data loss, calculating bandwidth based on 36
- date formats, supported 16
- diskStats command, collecting data 25
- downtime constraint, Secondary 44

F

- files
 - analysis results 38

G

- graphs, displaying analysis results 33

I

- installation wizard 12
- installing VRAdvisor
 - on Solaris 11
 - on Windows 12

J

- Japanese locales, supported 16

L

- layered volumes, collecting data on 30
- locales, supported 16

M

- maximum outage window 34
- minimum time for collecting data 15

N

- network bandwidth
 - choosing 40
- network bandwidth, required for synchronous replication 34

O

- operating systems, supported 11
- outage window, maximum 34
- outage, calculating SRL size from 36

P

- peak usage constraint 40
- period of synchronization 42
- permissible outage, calculating SRL size from 36
- pkgadd command 11
- pkgrm command 11
- platforms, supported 11
- preferences, modifying 37

R

- RAID-5 volumes not supported 30
- recalculating analysis results 34
- removing VRAdvisor
 - on Solaris 11
 - on Windows 13
- required bandwidth, synchronous replication 34
- required SRL size 34
- results file, location 38
- results graphs, described 33

S

- saving analysis results 38
- script, collecting data
 - on Solaris 19
- Secondary backup constraints 43
- Secondary backup interval
 - see Secondary outage 36
- Secondary downtime constraint 44
- slider bar, customizing values 37
- Solaris
 - installing VRAdvisor 11
 - uninstalling VRAdvisor 11
- SRL fillup 33
- SRL size
 - calculating 36
 - results of analysis 34
- supported locales 16
- supported operating systems 11
- synchronization period constraint 42
- synchronous replication, required bandwidth 34

U

- uninstalling VRAdvisor
 - on Solaris 11
 - on Windows 13
- UNIX
 - collecting data on 16
- usage constraint, peak 40

V

- value ranges, for slider bar 37
- VRAdvisor Wizard, collecting data
 - on Solaris 17
 - on Windows 22
- VRTSvradv.msi, installing 12

W

- what-if analysis
 - changing the value ranges 37
 - overview 35
 - performing 36
 - recording and viewing the results 38
- Windows
 - collecting data 22
 - installing VRAdvisor 12
 - uninstalling VRAdvisor 13