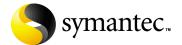# Veritas Volume Replicator Advisor User's Guide

Windows Server 2003, Windows Server 2008

5.1 Service Pack 1

symantec™

# Veritas Volume Replicator Advisor User's Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 5.1. Service Pack 1

Document version: 5.1.SP1.0

## Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

# Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's maintenance offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

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

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

## Contacting Technical Support

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

http://www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:

- Error messages and log files
- Troubleshooting that was performed before contacting Symantec
- Recent software configuration changes and network changes

## Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

http://customercare.symantec.com

## Customer service

Customer service information is available at the following URL:

http://customercare.symantec.com

Customer Service is available to assist with the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and maintenance contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

## Maintenance agreement resources

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

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

## Additional enterprise services

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

Enterprise services that are available include the following:

Symantec Early Warning Solutions    These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.

Managed Security Services    These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.

Consulting Services    Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.

Educational Services    Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

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

www.symantec.com

Select your country or language from the site index.

# Contents

# Introducing Veritas 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*.

## 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 17.

## 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 32.

# 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 2003: Standard Edition, Web Edition, Enterprise Edition, Datacenter Edition
- Windows 2008

## 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

**Note:** This section gives instructions on installing VRAdvisor on Windows. 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. 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 `common` directory under the `pkgs` directory on the CD.

3   Run the `vrtsvradv.msi` from the `common` 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
    ■   To choose another location for installing `VRAdvisor`, click **Change**.
    ■   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.**
    ■   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 2003 or Windows XP, a reboot is not required to enable disk
    performance counters.

## Uninstalling VRAdvisor on Windows

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**

■ Symantec recommends that you collect the sample data using the volumes that are part of the VVR configuration you are planning to set up.

■ To collect a representative sample of data, it is recommended 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.

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 18

■ "Collecting sample data on Windows" on page 23

# 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

■   When using Veritas Volume Manager (VxVM) volumes, make sure that you import the disk group containing the required volumes onto your system.

## 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   On the Welcome page, select **Data Collection**, and click **Next**.

5   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** | 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. |
| | Select the required volumes individually, or click **Select All** to select all of the available volumes in the selected disk group. |

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, and *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 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:

　■ For AIX, use lvmstat to collect data for all logical volumes.

　■ For HP-UX and Linux use sar to collect data for all disks.

　■ For Solaris, use 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 the 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 run the following command:**

◆   # **sh vra_data_collection.sh -h**

## Examples

### Example 1
To collect data every minute for 30 hours:

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

### Example 2
To 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
To collect data for VxVM disk groups dg1 and dg2:

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

### Example 4
To collect data for VxVM volumes vol1, vol2, and vol3 in disk group dg1:

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

### Example 5
To 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 sample data on Windows

VRAdvisor can be used to collect and analyze a sample 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 data using one of the following methods:

- Collecting sample data using the VRAdvisor wizard
- Collecting the sample data using the diskStats command

### Prerequisite

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

## Collecting sample data using the VRAdvisor wizard

**To collect data using the VRAdvisor wizard**

1   To launch the VRAdvisor wizard on Windows, select **Start > All Programs > Symantec > Volume Replicator Advisor > VRAdvisor Wizard**.

2   On the Welcome page, select **Data Collection** and click **Next**.
The Data Collection page appears.

**Note:** On Windows, only the diskStats command is 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**. VRAdvisor continues to collect data for the specified duration, although the wizard window closes.
The data collection wizard displays an 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 the sample 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:

    ```
    diskStats  [-i interval [-c count]] \
        <drive 1> [[drive 2][drive 3]... ]
    ```
    The command will display the output on the console.

---

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

---

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 data collection completes, the file *filename* contains the sample data in `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 data that you have collected. VRAdvisor analyzes the sample 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 (Storage Replicator Log) 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 37.

VRAdvisor enables you to analyze data collected on any of the supported platforms described in "Collecting the sample of data" on page 17. However, to analyze the data, you must install and use VRAdvisor on either Solaris or Windows system.

### Prerequisites

- All the files to be analyzed must be present in a single directory.
- The sample 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.

    For Solaris:

    ■   Change directory as follows:

        # **cd /opt/VRTSvradv/bin**

    ■   Launch the VRAdvisor wizard, using the following command:

        # ./**vradvgui**

    For Windows:

    Choose **Start** > **Programs** > **VERITAS** > **Volume Replicator Advisor** > **VRAdvisor Wizard**.

3   On the Welcome page, select **Analysis**, and then click **Next**.

4   On the Directory Specification page, enter the name of the directory containing the data files to be analyzed. All files to be analyzed must be present 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.

---

    ■   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 result. You can also change some parameters and recalculate the result. You can perform the following actions:

■ 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 is displayed by default.

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

### Analysis graph

The Analysis Graph section shows the following information:

■ The top graph shows the SRL (Storage Replicator Log) fillup in megabytes (MB) on the y-axis. The fillup rate is shown both with specified outages and

no outages. 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 a worst case scenario.

**Note:** If SRL fillup value in the graph steadily increases upto maximum during the last data write duration, it indicates that you do not have sufficient network bandwidth for the amount of data writes contained in the sample data.

■   The bar graph shows the value 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 specified for network bandwidth and 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:** Symantec recommends that you add a 10-20 percent buffer to the values calculated by VRAdvisor when setting up the 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 as 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 (Recovery Point Objective) and the RTO (Recovery Time Objective)
that can be achieved after a disaster. Use this option to calculate the
bandwidth required for a Primary and Secondary site and between a
Bunker and Secondary based on the desired RPO and RTO.
Available parameters for this option are % Disk Writes and RTO. The
**Have Bunker** check box indicates that 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 34.

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 on 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 values that you specified during 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 at the `VERITAS/Volume Replicator Advisor/results` subdirectory.

For Solaris, the file is located at 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 Storage Replicator Log (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 formulas described in this section to determine the appropriate SRL size.

---

**Note:** The terms Replicator Log and Storage Replicator Log (SRL) mean the same.

---

## 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 a SRL must be at least 110 MB. If size specified for SRL is less than 110 MB, VVR displays an error message that prompts to specify a value that is equal to or greater than 110 MB.

---

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

- The maximum expected downtime for Secondary nodes

- The maximum expected downtime for a 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 data over the 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 Automatic Synchronization option is used to synchronize the Secondary, the above-mentioned step is not a concern.

To perform Secondary backups inorder to avoid complete resynchronization in case of Secondary data volume failure, the following information is required:

- 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 certain period 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. However, for synchronous RLINKs this does not apply as a shortfall in network capacity would cause each application write to be delayed. This in turn causes the application to run more slowly.

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 ($BW_N$).

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

3   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 39, 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.

Several factors can reduce the maximum size to which the SRL can fill up during the peak usage period. The factors that need to be considered 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 |
| 9 | 10 | 15 | 5 | 10 | 16 |
| 10 | 11 | 15 | 5 | 10 | 26 |

**Table A-1**      Example calculation of SRL size required to support peak usage
period(Continued)

| Hour Starting | Hour Ending | Application (GB/hour) | Network (GB/hour) | SRL Usage (GB) | Cumulative SRL Size (GB) |
|---------------|-------------|----------------------|-------------------|----------------|--------------------------|
| 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 |

## 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.*

If SRL overflows during the synchronization period when the application is running and data is getting accumulated in the SRL, then the synchronization process must be restarted. To ensure that the SRL does not overflow during such periods, it is necessary to appropriately size the SRL so that it can hold as much data as the application writes. After replication is started, 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 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 40, 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. If a failure occurs, then the Secondary data volumes are restored from this local copy and replication proceeds from the checkpoint. Data is replayed 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.

■ Minimum SRL size.
  You need to consider an application's write rate and frequency of Secondary backups in order calculate the minimum 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 scenario the amount of data the application could generate over this time period.

---

**Note:** Even if only one volume fails, all other volumes need to 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 and hence, 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 would be repaired within a stipulated period. However, if the repair is not completed within the guaranteed period, then SRL may overflow. Hence, it is recommended that a safety margin should always be added to any such arrived 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 scenario, 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 likely 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 as after a disconnect they act as asynchronous RLINKs until caught up.

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% |
| 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 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 maximum 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 GB/hour x 12 hours+ 1/4 GB/hour x 12 = 15 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