

Technote: Adaptive Synchronous Replication in InfoScale

7.4

Last updated: 2019-04-01

Legal Notice

Copyright © 2019 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third-party ("Third-Party Programs"). Some of the Third-Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third-party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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 Veritas Technologies LLC and its licensors, if any.

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

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

<https://sort.veritas.com/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

infoscaledocs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Setting up and configuring adaptive synchronous replication

This document includes the following topics:

- [Understanding replication settings](#)
- [Modes of VVR replication](#)
- [About adaptive synchronous replication](#)
- [Configuring adaptive synchronous replication](#)
- [Limitation](#)

Understanding replication settings

The VVR replication settings determine the replication behavior between the Primary RVG and a specific Secondary RVG. VVR behaves differently based on the settings for mode of replication, SRL overflow protection, and latency protection, depending on whether the Secondary is connected or disconnected. To use the replication settings effectively, it is important to understand how each attribute affects the replication when the Primary and Secondary are connected and disconnected.

VVR enables you to set up the replication mode, latency protection, and SRL protection using the replication attributes. Each attribute setting can affect replication and must be set up with care.

For details on VVR, see the *Veritas InfoScale Replication Administrator's Guide*.

Modes of VVR replication

VVR replicates in synchronous and asynchronous modes. Before you choose synchronous or asynchronous mode, you must understand the effects of this choice on the replication process and the application performance.

- **Synchronous replication:** Choose the synchronous mode if, in the case of a disaster, the Secondary must reflect all writes that have successfully completed on the Primary.
- **Asynchronous replication:** Choose the asynchronous mode if some delay is acceptable before the Secondary data volumes are updated with the writes completed on the Primary volumes.
- **Adaptive synchronous replication:** Choose this replication mode if you want the replication mode to automatically switch from synchronous to asynchronous and vice versa based on cross-site network latency.

Note: The Adaptive synchronous replication feature is available only when you install 7.4.0.1500 Patch for InfoScale 7.4.

You can configure the mode of replication by setting the value of the `synchronous` attribute to `override`, `off`, or `fail`.

Table 1-1 Mode of replication and the state of the RLINK

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

About adaptive synchronous replication

When the `synchronous` attribute is set to `override`, synchronous mode switches to asynchronous mode during a temporary network outage. After the outage passes, the Secondary catches up and replication reverts to synchronous.

The `override` option allows VVR to continue receiving writes from the application in asynchronous mode even when RLINK is disconnected. However, in case of high network latency, replication continues to run in synchronous mode with degraded application performance.

The adaptive synchronous mode in VVR is an enhancement to the existing synchronous override mode. In the adaptive synchronous mode, replication switches from synchronous to asynchronous based on cross-site network latency. This allows replication to take place in synchronous mode when network conditions are good, and automatically switch to asynchronous mode when there is an increase in cross-site network latency. You can also set alerts to notify you if the system undergoes network deterioration. The alerts are logged in the system log file and the `/etc/vx/log/vxloggerd.log` file.

You must define the following attributes to enable the adaptive synchronous replication.

Table 1-2 Required attributes for adaptive synchronous replication

Attribute	Description
<code>synchronous</code>	Set the <code>synchronous</code> attribute of the RLINK to <code>override</code> . By default, VVR sets the <code>synchronous</code> attribute to <code>off</code> . Setting the <code>synchronous</code> attribute to <code>override</code> puts the RLINK in synchronous mode and specifies override behavior if the RLINK is disconnected.
<code>iotimeout</code>	Define the <code>iotimeout</code> attribute to specify the time interval, in microseconds (μ s), for which the Primary waits for an acknowledgment from the Secondary, before considering an update to be timed out. The <code>iotimeout</code> attribute must be defined by the user and is required by VVR to replicate to a Secondary in the adaptive synchronous mode.

Additionally, you can define the following attributes to modify the override behavior of the adaptive synchronous replication.

Table 1-3 Optional attributes for adaptive synchronous replication

Attribute	Description	Default Value
<code>threshold</code>	Maximum permissible value, in percentage, of the updates that were sent to the Secondary and were timed out within a user-defined time interval. VVR uses the threshold value to determine when to switch from synchronous to asynchronous mode.	30%

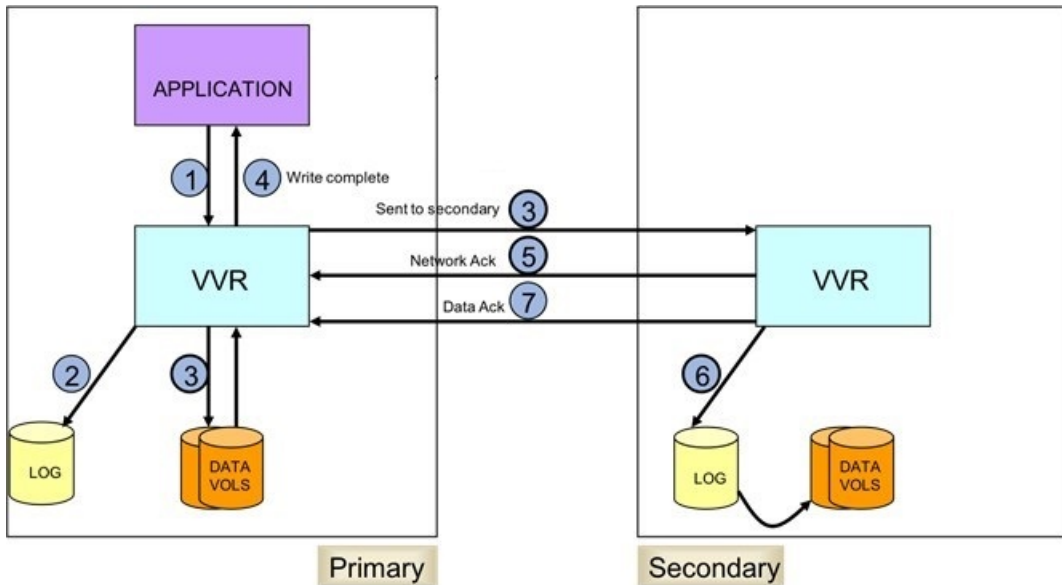
Table 1-3 Optional attributes for adaptive synchronous replication
(continued)

Attribute	Description	Default Value
interval	Time interval, in seconds, that is used for calculating the percentage of timed out updates.	60 seconds
switch_timeout	Minimum time interval, in seconds, for which the system must remain in asynchronous mode, before switching back to synchronous mode.	300 seconds
threshold_monitor	<p>Threshold value, in percentage, of timed out updates. This attribute is used to notify the user in advance of a potential degradation in the network.</p> <p>When the percentage of timed out updates exceeds this value, it is considered as an early indication of network deterioration, and an email alert is sent to the user.</p>	5%
async_backlog_threshold	If replication is running in the asynchronous mode for period longer than the value specified in this attribute, an email notification is sent to the user.	300 seconds
autowitch	<p>This attribute specifies whether the system automatically switches between synchronous and asynchronous replication mode based on the threshold value, that is, the network conditions. By default, this value is set to off.</p> <p>Set the autowitch value to on, if you want the system to switch to the asynchronous mode when the network is degraded and switch back to synchronous mode when network conditions improve.</p> <p>Note: Email notifications for <code>threshold_monitor</code> and <code>async_backlog_threshold</code> attributes are sent by the system irrespective of whether <code>autowitch</code> is set to on or off.</p>	off

How adaptive synchronous replication works

Note that in the following figure, the repeating numbers denote parallel activities.

Figure 1-1 Flow of timed-out IO in adaptive sync replication



In the adaptive synchronous mode of replication, VVR processes an incoming write by performing the following steps in the following order:

- VVR receives a write on the Primary.
- VVR processes the incoming write in synchronous replication mode:
 - Writes it to the Primary SRL.
 - VVR writes to the data volumes on the Primary. At the same time, VVR sends the write to the Secondary hosts and waits for the synchronous network acknowledgments from the Secondary hosts.
 - If Primary does not receive an acknowledgment within the set `iotimeout` period, it considers the update to be timed out and sends an advance acknowledgment to the application.
 - If the Percentage of such timed-out updates exceeds the set threshold value, which indicates that the network is degraded, the mode of replication changes to asynchronous.
- VVR continues to process the incoming write in asynchronous replication mode.

- If the network conditions improve during this period, the system automatically switches the replication mode and starts operating in synchronous mode.

Alerts and email notifications

For better monitoring, VVR uses VxLogger to log the alerts messages for different conditions. These alerts are logged in the system log file and the `/etc/vx/log/vxloggerd.log` file. These alerts are also sent to the configured users as email notifications.

Alerts are logged, and email notifications are sent in the following events:

Event	Alert Message
Replication mode changes from synchronous to asynchronous mode.	ADSYNC_ALERT VxVM VVR vxio V-5-1-0 Switching adaptive replica <rlink> to asynchronous mode. IO timeout threshold reached
Replication mode changes from asynchronous to synchronous mode.	ADSYNC_ALERT VxVM VVR vxio V-5-1-0 Switching adaptive replica <rlink> to synchronous mode
<code>Threshold_Monitor</code> percentage value exceeds the set limit that indicates a potential service deterioration.	ADSYNC_ALERT VxVM VVR vxio V-5-1-0 Around 20% IOs timed out on replica <rlink> in interval of 60 seconds
Replication takes place in asynchronous mode for more than a defined time period.	ADSYNC_ALERT VxVM VVR vxio V-5-1-0 Replica <rlink> is running in asynchronous mode for 300 seconds

Note: The email notifications are processed by the `vxloggerd` daemon. To receive alerts through email, you must ensure that the `vxloggerd` daemon is running. If you erroneously stop this daemon, it does not start automatically, and must be started manually.

By default, the alert email notifications are sent to the root user. The root user may change in case of log owner switch.

We can also send the email notification to the specified recipients. To do this, add the comma separated list of recipients in the `/etc/vx/mail_recipients` configuration file. For example, `user1@example.com, user2@example.com`, and so on. Ensure that you maintain the same recipient list in the configuration file on each node of the cluster.

Configuring adaptive synchronous replication

- 1 Create the Primary RVG.

```
# vradmin -g diskgroup createpri local_rvgname data_volume  
log_volume
```

Where *data_volume* is name of the data volume and *log_volume* is the name of the SRL log volume.

- 2 Create the Secondary RVG.

```
# vradmin -g diskgroup addsec local_rvgname <pri_hostname>  
<sec_hostname>
```

The argument *pri_hostname* and *sec_hostname* are the name of the Primary and Secondary hosts respectively.

- 3 To set the replication to adaptive synchronous mode, set the synchronous attribute to override.

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

- 4 To set the adaptive synchronous replication parameters, run the following command:

```
# vradmin -g diskgroup set local_rvgname iotimeout=1000  
threshold=30 interval=20 switch_timeout=300
```

- 5 To start replication using automatic synchronization, run the following command:

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

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

The argument *pri_hostname* and *sec_hostname* are the name of the Primary and Secondary hosts respectively. If the RDS contains only one Secondary, the *sec_hostname* is optional.

Limitation

The adaptive synchronous replication mode does not support replication over multiple secondary sites.