

Arcserve® Replication and High Availability

Oracle Server Operation Guide for Windows r16.5

arcserve®

Pre-release Document, only for reference

This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by Arcserve at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of Arcserve. This Documentation is confidential and proprietary information of Arcserve and may not be disclosed by you or used for any purpose other than as may be permitted in (i) a separate agreement between you and Arcserve governing your use of the Arcserve software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and Arcserve.

Notwithstanding the foregoing, if you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all Arcserve copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to Arcserve that all copies and partial copies of the Documentation have been returned to Arcserve or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, ARCSERVE PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL ARCSERVE BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF ARCSERVE IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is Arcserve.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

© 2015 Arcserve, including its affiliates and subsidiaries. All rights reserved. Any third party trademarks or copyrights are the property of their respective owners.

Arcserve Product References

This document references the following Arcserve products:

- Arcserve® Replication
- Arcserve® High Availability (HA)
- Arcserve® Assured Recovery®
- Arcserve® Content Distribution

Contact Arcserve

The Arcserve Support team offers a rich set of resources for resolving your technical issues and provides easy access to important product information.

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

With Arcserve Support:

- You can get in direct touch with the same library of information that is shared internally by our Arcserve Support experts. This site provides you with access to our knowledge-base (KB) documents. From here you easily search for and find the product-related KB articles which contain field-tested solutions for many top issues and common problems.
- You can use our Live Chat link to instantly launch a real-time conversation between you and the Arcserve Support team. With Live Chat, you can get immediate answers to your concerns and questions, while still maintaining access to the product.
- You can participate in the Arcserve Global User Community to ask and answer questions, share tips and tricks, discuss best practices and participate in conversations with your peers.
- You can open a support ticket. By opening a support ticket online, you can expect a callback from one of our experts in the product area you are inquiring about.

You can access other helpful resources appropriate for your Arcserve product.

Documentation Changes

The following documentation updates have been made since the last release of this documentation:

- Updated to include user feedback, enhancements, corrections, and other minor changes to help improve the usability and understanding of the product or the documentation itself.

Contents

Chapter 1: Introduction	7
About This Guide	7
Related Documentation	7
Server Requirements.....	8
Log On Account Conditions.....	8
Servers Operating in a Workgroup.....	9
Chapter 2: Redirection Methods	11
How Redirection Works	11
DNS Redirection	11
Move IP Redirection.....	12
Scripts Redirection	12
Add IP on Master Server	12
Chapter 3: Creating and Using Scenarios	13
Considerations.....	14
Create Oracle HA and DR Scenario.....	15
Scenario Properties	21
Run a Scenario.....	23
Stop a Scenario.....	25
View a Report.....	25
Chapter 4: Switching Over and Switching Back	29
How Switchover and Switchback Work.....	30
Initiate Switchover	31
Initiate Switchback	33
Switchover Considerations.....	36
Chapter 5: Recovering Data	37
The Data Recovery Process	37
How to Recover Oracle Data	37
Setting Bookmarks	38
Data Rewind	39
Recover Active Server	41

Appendix A: Additional Information and Tips	43
Pool Directory Settings	43
Manually Recover a Failed Server-Move IP Address	44
Troubleshooting Oracle Databases	44
Cannot Connect Oracle Enterprise Manager After Switchover	45
Index	47

Chapter 1: Introduction

Arcserve RHA (RHA) is a solution based on asynchronous real-time replication and automated application switchover and switchback to provide cost-effective business continuity for Oracle and other application servers on both 32-bit and 64-bit Windows servers.

Arcserve RHA lets you replicate data to a local or remote server, making it possible to recover that data due to server crash or site disaster. You may switch your users to the replica server manually, or automatically, if you licensed High Availability. This Guide presents both Replication and High Availability concepts and procedures.

The procedures covered in this Guide should be followed as is. Customize the steps only if:

- You are familiar with Arcserve RHA and fully understand the potential impact of any changes you make.
- You have fully tested the steps in a lab environment before implementing in a production environment.

This section contains the following topics:

[About This Guide](#) (see page 7)

[Related Documentation](#) (see page 7)

[Server Requirements](#) (see page 8)

About This Guide

This document describes how you can implement a Arcserve RHA solution for Oracle. It is essential that you have the appropriate resources and permissions to carry out each task.

Related Documentation

Use this Guide with the *Arcserve RHA Installation Guide* and the *Arcserve RHA Administration Guide*.

Server Requirements

To implement Arcserve RHA, refer to the appropriate list of requirements, depending on the server type you selected. These components are licensed separately. If you do not have the license required to access support for a specific server type, contact Technical Support.

(In the Active Directory environment) Both Master and Replica servers should reside in the same Active Directory forest and be members of the same domain or trusted domains.

For Oracle Server configuration:

- Two servers running supported Windows Server with the same level of service packs and hot fixes installed.
Note: For a complete list of supported operating systems and applications, see the Arcserve RHA Release Notes.
- Install the same Oracle version on both servers, including service packs and hot fixes.
- The Oracle SID on the Master must match the Replica.
- Ensure the path to the ORACLE_HOME directory and the path to the database files are identical on both Master and Replica.
- Statically assign all IP addresses (DHCP-assigned IPs are not supported on either server).
- The protected server cannot be a domain controller or DNS server.
- To minimize network traffic, Oracle temporary tablespace names and path are excluded from replication. Make sure the Oracle database on the Replica server is configured with the same temporary tablespace names and path used on the Master server.
- For a 64-bit Windows server, you must have a 64-bit Oracle Database Instant client. For more information about the Oracle Database Instant client, visit the Oracle website.

Notes:

- Arcserve HA for Oracle does not support Microsoft cluster servers
- Arcserve HA for Oracle does not support Oracle RAC

Log On Account Conditions

The Arcserve RHA Engine service must satisfy certain account conditions for successful communication with other components. If these requirements are not met, scenarios may not run. If you lack the permissions required, contact your local IS team.

- It is a member of the Domain Admins group. If the Domain Admins group is not a member of the built-in domain local group Administrators you must use an account that is.
- It is a member of the local computer Administrators Group. If the Domain Admins group is not a member, add the account manually.
- For servers in a workgroup, use the Local System account. If you have used Redirect DNS redirection method in a HA scenario, then use the local administrator account.

Servers Operating in a Workgroup

For servers in a workgroup, set the Arcserve RHA Engine service account to Local System (unless you have added any custom scripts that require higher level permissions). Oracle servers in a workgroup can use Redirect DNS only with DNS servers that allow non-secure updates. You can use Move IP, and custom redirection scripts normally, but the Switch Computer Name method is not supported.

Chapter 2: Redirection Methods

This section contains the following topics:

[How Redirection Works](#) (see page 11)

[DNS Redirection](#) (see page 11)

[Move IP Redirection](#) (see page 12)

[Scripts Redirection](#) (see page 12)

[Add IP on Master Server](#) (see page 12)

How Redirection Works

Each server type supported by Arcserve RHA can be configured to use one or more redirection methods. You should enable redirection methods based on your environment and business needs. The following sections describe the supported redirection methods for Oracle.

Note: For Microsoft Exchange Server 2010, by default, only Move IP is available for the redirection method. In addition, Exchange Server 2010 HA scenario works fine even if all redirection methods are set to off.

DNS Redirection

Important! This is the recommended redirection method for Oracle HA scenarios. During scenario creation, this setting is On by default.

DNS Redirection changes the Master server DNS to the Replica server IP address. Upon failure of the Master, the Replica server modifies the appropriate DNS record so that references to the Master server resolve to the Replica's IP address rather than the Master's IP address. This redirection method requires no network reconfiguration and works in LAN and WAN network configurations.

DNS redirection works only with A (host) type records and cannot update CNAME (Alias) records directly. However, if the CNAME record points to the modified A record, it is indirectly redirected.

Using the record that has the Master server's name is the default, however you can configure Arcserve HA to redirect any DNS A (host) record via the *Master's name in DNS* setting in the switchover properties tab.

Note: If you use Switch Hostname as the redirection method in UNIX/Linux scenarios, you should also set DNS redirection for those scenarios.

Move IP Redirection

Move IP redirection involves moving the Master server IP address to the Replica server.

This redirection method is preferred for Virtual Machine scenarios and is usable only in a LAN configuration in which the Master and Replica servers reside in the same network segment. In this configuration, switchover of the Master server causes the Replica to take over one or more of the IP addresses assigned to the Master server.

Important! Use this method only when both servers are on the same IP subnet.

When using Move IP as the redirection method, you must first add IP addresses to the Master host. For more information, refer to the topic, Add IP on the Master Server.

Scripts Redirection

Arcserve RHA can trigger custom scripts or batch files to perform the user redirection or any additional steps not covered by built-in methods. If the above methods are not appropriate or do not fully meet all requirements, see the *Arcserve RHA Administration Guide* for details on scripted redirection methods.

Add IP on Master Server

To add IP addresses on the Master server, refer to your operating system documentation.

Chapter 3: Creating and Using Scenarios

The following topics describe how to create and configure a replication scenario.

This section contains the following topics:

[Considerations](#) (see page 14)

[Create Oracle HA and DR Scenario](#) (see page 15)

[Scenario Properties](#) (see page 21)

[Run a Scenario](#) (see page 23)

[Stop a Scenario](#) (see page 25)

[View a Report](#) (see page 25)

Considerations

Review the following consideration before you create a scenario:

You must have a 64-bit Oracle Database Instant client on a 64-bit Windows server. If you have a 32-bit Oracle client on a 64-bit Windows server, perform the following steps:

1. Download and install the 64-bit Oracle Database Instant client. For more information about the Oracle Database Instant client, visit the Oracle website.
2. Set the OCI library file path to the 64-bit OCI.DLL file while creating the scenario using the scenario wizard.
3. Edit the Net Service Name to include the server name or IP address in the following format:

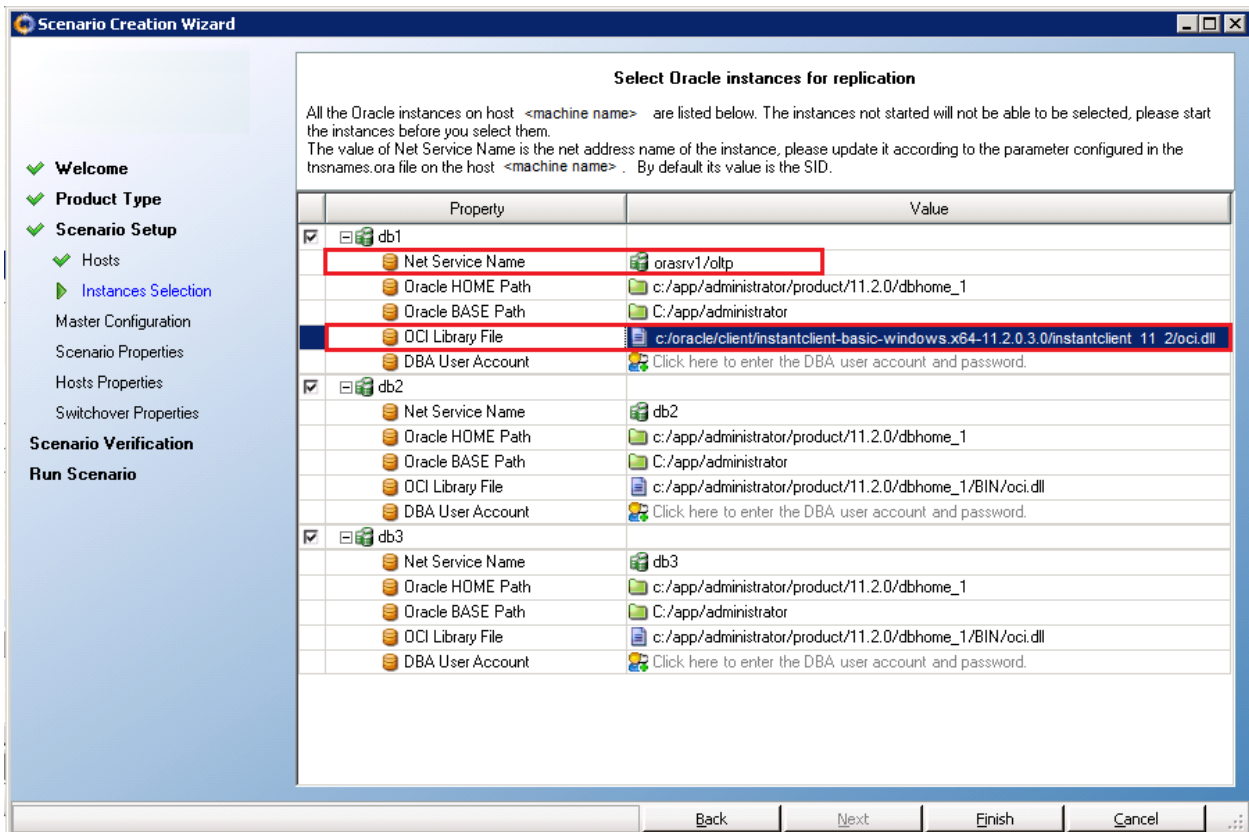
<Database Server Name>/<Database Name>

or

<Database Server IP address>/<Database Name>

Example:

192.xx.xx.10/DB1



Create Oracle HA and DR Scenario

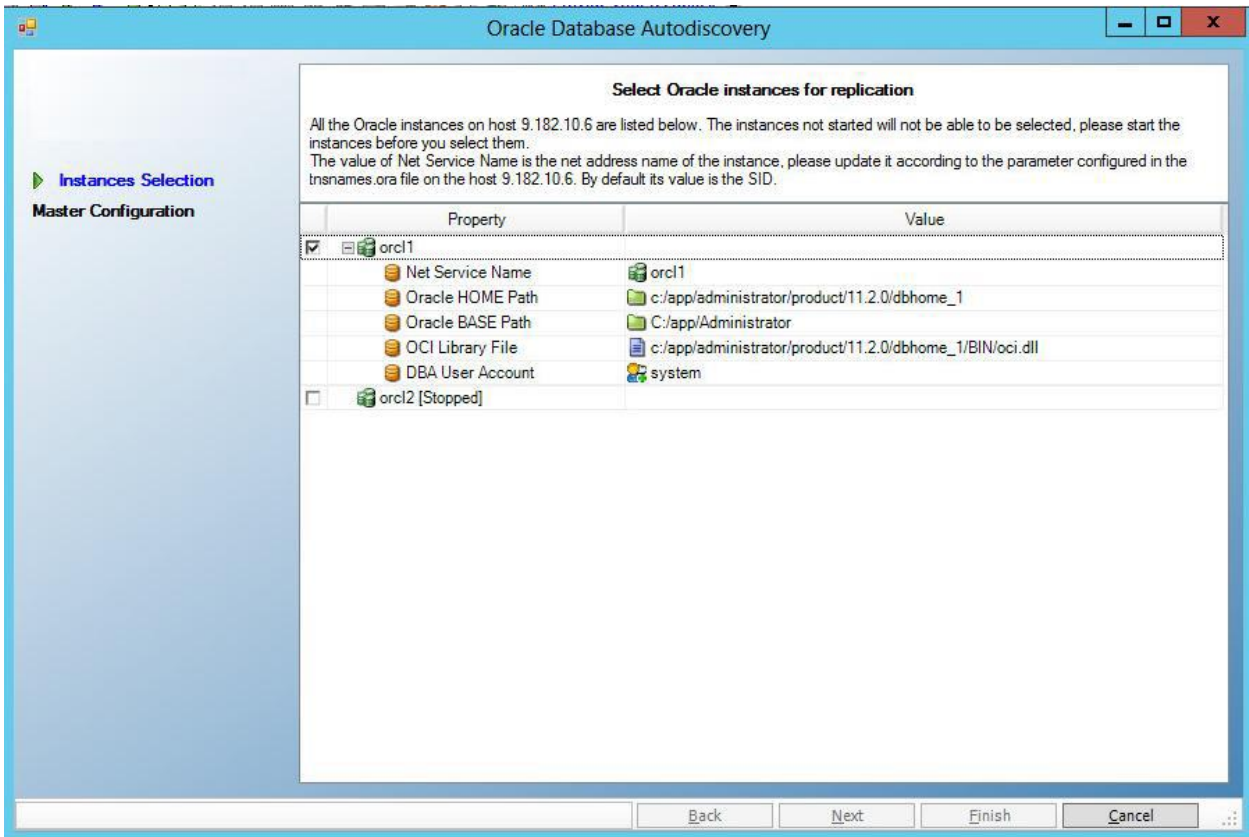
Protecting Oracle databases using Arcserve RHA requires scenarios that identify the Master and Replica servers, as well as properties that control scenario behavior. Auto-discovery is used to protect multiple Oracle instance in one scenario. RHA monitors all services for each Oracle instance and failure of a service triggers switchover.

RHA retrieves multiple Oracle instances from the registry. RHA also gets the ORACLE_HOME and ORACLE_HOME/ORACLE_BASE path from the Windows registry.

Follow these steps:


1. Start Arcserve RHA Manager. From the Scenario menu, choose New or click the New Scenario button.
2. From the Welcome dialog, choose Create a New Scenario; choose a scenario group and click Next.
3. From the Select Scenario and Product Type dialog, select Oracle, High Availability Scenario (HA), the desired Tasks on Replica, if any, and then click Next.

- From the Master and Replica Hosts dialog, type a Scenario Name, enter the Hostname or IP Address and Port number for both Master and Replica servers. Click Next.
- From the Select Oracle Instances for replication screen, select one or more Oracle DB instances to replicate the scenario. On this screen, auto-discovery displays all Oracle instances. These instances are selected automatically. You can reconfigure the parameters including Net Service Name, Oracle Base Path, Oracle Home Path, OCI library file, and the user account for DB connection for each Oracle instance. You can then select the instances that you want to replicate.



6. Click DBA User Account.

The Database Connection dialog opens.



The screenshot shows a "Database Connection" dialog box. The title bar includes a help icon (question mark) and a close icon (X). The main content area features a database icon and the instruction "Please provide login parameters:". Three input fields are present: "Database Name" containing "orcl1", "DBA Username" containing "sys as sysdba", and "Password" containing "*****". At the bottom, there are three buttons: "OK", "Cancel", and "Apply to All".

7. Enter the database name, DBA username, and password and click OK.

Note: To apply credentials for all selected instance, click Apply to All.

Enter credentials for an account with administrative privileges, such as "sys as sysdba", in the database so that Arcserve RHA can query the Oracle configuration on the Master server.

For the database name, use the database "connection string" configured in the Oracle client configuration file tnsnames.ora on the Master server.

Note: If you are using a 32-bit Oracle version on a 64-bit operating system, enter the global service name in the Database Name field, in the following format:

IP <:port>/global_service_name. Omit the <:port> if the default port (1521) is used.

Example:

192.xx.xx.10/db1.test.com

Enter the DBA Username and Password as usual and click OK.

Also make sure that OCI Library File on the Select Oracle Instances for replication screen is updated to point to the 64-bit Oracle client OCI library file.

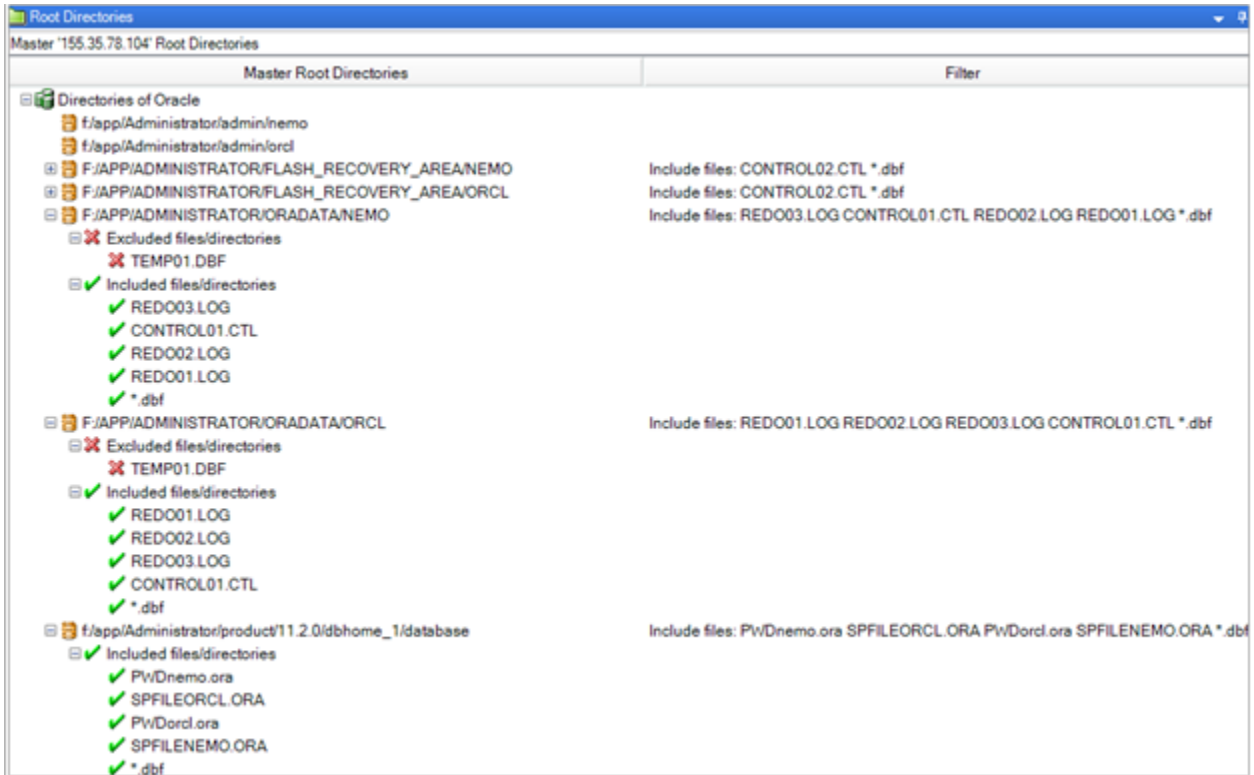
The Tablespaces for Replication dialog opens.

8. Review the list of results auto-discovered for each instance on the Master Server. Clear and select databases for replication, as desired, enable the "Replicate new user created files in listed root directories" option and click Next.

Selecting this option automatically replicates tablespaces to Oracle databases as you add them. If this option is not selected, Arcserve RHA replicates only the databases that are selected at configuration time.

	Tablespaces	Full Path
<input checked="" type="checkbox"/>	Oracle	
<input checked="" type="checkbox"/>	orcl	
<input checked="" type="checkbox"/>	Common	
<input checked="" type="checkbox"/>	Password File	f:/app/Administrator/product/11.2.0/dbhome_1/database/PWDorcl.ora
<input checked="" type="checkbox"/>	Parameter File	F:/APP/ADMINISTRATOR/PRODUCT/11.2.0/DBHOME_1/DATABASE/SPFILEORCL.ORA
<input checked="" type="checkbox"/>	Admin	f:/app/Administrator/admin/orcl
<input checked="" type="checkbox"/>	orcl	
<input checked="" type="checkbox"/>	Control files	
<input checked="" type="checkbox"/>	File	F:\APP\ADMINISTRATOR\ORADATA\ORCL\CONTROL01.CTL
<input checked="" type="checkbox"/>	File	F:\APP\ADMINISTRATOR\FLASH_RECOVERY_AREA\ORCL\CONTROL02.CTL
<input checked="" type="checkbox"/>	Log files	
<input checked="" type="checkbox"/>	File	F:\APP\ADMINISTRATOR\ORADATA\ORCL\REDO03.LOG
<input checked="" type="checkbox"/>	File	F:\APP\ADMINISTRATOR\ORADATA\ORCL\REDO02.LOG
<input checked="" type="checkbox"/>	File	F:\APP\ADMINISTRATOR\ORADATA\ORCL\REDO01.LOG
<input checked="" type="checkbox"/>	Table Spaces	
<input checked="" type="checkbox"/>	SYSTEM	F:\APP\ADMINISTRATOR\ORADATA\ORCL\SYSTEM01.DBF
<input checked="" type="checkbox"/>	SYSAUX	F:\APP\ADMINISTRATOR\ORADATA\ORCL\SYSAUX01.DBF
<input checked="" type="checkbox"/>	UNDOTBS1	F:\APP\ADMINISTRATOR\ORADATA\ORCL\UNDOTBS01.DBF
<input checked="" type="checkbox"/>	USERS	F:\APP\ADMINISTRATOR\ORADATA\ORCL\USERS01.DBF
<input type="checkbox"/>	TEMP	F:\APP\ADMINISTRATOR\ORADATA\ORCL\TEMP01.DBF
<input checked="" type="checkbox"/>	EXAMPLE	F:\APP\ADMINISTRATOR\ORADATA\ORCL\EXAMPLE01.DBF
<input checked="" type="checkbox"/>	nemo	
<input checked="" type="checkbox"/>	Common	
<input checked="" type="checkbox"/>	Password File	f:/app/Administrator/product/11.2.0/dbhome_1/database/PWDnemo.ora
<input checked="" type="checkbox"/>	Parameter File	F:/APP/ADMINISTRATOR/PRODUCT/11.2.0/DBHOME_1/DATABASE/SPFILENEMO.ORA
<input checked="" type="checkbox"/>	Admin	f:/app/Administrator/admin/nemo
<input checked="" type="checkbox"/>	nemo	
<input checked="" type="checkbox"/>	Control files	
<input checked="" type="checkbox"/>	File	F:\APP\ADMINISTRATOR\ORADATA\NEMO\CONTROL01.CTL
<input checked="" type="checkbox"/>	File	F:\APP\ADMINISTRATOR\FLASH_RECOVERY_AREA\NEMO\CONTROL02.CTL
<input checked="" type="checkbox"/>	Log files	

9. From the Root Directories dialog, review the list of folders to be replicated and click Next.



10. From the Scenario Properties dialog, configure the settings that control the entire scenario. For more information, see the Arcserve RHA Administration Guide. For Oracle scenarios, we recommend enabling the Replicate NTFS ACL option, under Replication, Optional Settings. Click Next.
11. From the Master and Replica Properties dialog, configure the desired settings and then click Next.
12. Wait for the Switchover Properties dialog to appear. Configure the redirection settings as described in IP Redirection, and then click Next.

Note: You can only use Move IP and Redirect DNS for an Oracle scenario.
13. From the Switchover Initiation dialog, configure the desired Initiation options and click Next. We recommend setting only one of these to Automatic.

Note: If you selected Assured Recovery in the Tasks on Replica options, that dialog appears now. For more information, see the Arcserve RHA Administration Guide.

14. From the Scenario Verification dialog, ensure that no errors or warnings appear. While warnings do not prevent a scenario from running as errors do, both should be corrected before you proceed. Once successfully verified, click Next. The pre-run checks for HA verify master and replica settings, such as Oracle home, for each instance.

Note: If the Oracle home path on the master is changed, reconfigure the parameters of Oracle instances and perform the auto-discovery again.

15. Click Run Now to start data synchronization or click Finish to run the scenario later.

Scenario Properties

If you wish to change a scenario configured through the Wizard or configure additional settings, you can use the Properties pane to modify the scenario.

The Properties pane and its tabs are context-sensitive and change whenever you select a different node from a scenario folder. You must stop a scenario before configuring its properties. Certain values cannot be modified once set; they are noted.

For full details on configuring scenario properties and their descriptions, see the Arcserve RHA Administration Guide.

Properties are organized into tabs on the Arcserve RHA Manager Framework pane. The tabs displayed are based upon server type, Arcserve RHA solution, and scenario status.

Settings on the Root Directories tab

Select a Master Server from the Scenario Pane. Double-click its Directories folder to add or remove Master Root Directories. Select or clear checkboxes next to folders, as desired, to include or exclude them. You may also edit directory names.

Select a Replica Server from the Scenario Pane. For each Master Root directory, you must specify a Replica Root directory. Double-click the Directories folder for the Replica server. Select or clear checkboxes next to folders, as desired, to hold the corresponding Master directory.

Directories with Oracle database files are automatically included by Oracle auto-discovery in the scenario at the scenario configuration step (see the topic, Create Oracle HA Scenarios). There is no need to add these directories manually. If you want to include additional directories in the scenario, you can add them using the Root Directories tab. We strongly recommend not including directories with Oracle binaries, DLLs and other Oracle instance directories that are not part of the database.

Settings on the Properties Tab

Scenario Properties

These settings establish default behavior for the entire scenario.

- General properties -- cannot be changed once created
- Replication properties -- choose the replication mode (Online or Scheduled), synchronization values (File or Block, Ignore Files of Same Size/Type) and optional settings (Replicate NTFS Compress Attribute, Replicate NTFS ACL, Synchronize Windows Shares, Prevent Automatic Re-sync upon Error)
- Event notification properties -- specify a script to run or choose email notification
- Report Handling -- specify report settings, email distribution or script execution

Master and Replica Properties

These settings establish server properties on both Master and Replica. Some settings vary by server type.

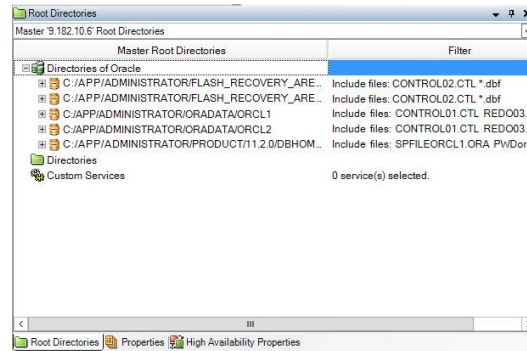
- Host connection properties -- Enter the IP address, Port number and Fully Qualified Name of the Master and Replica
- Replication properties -- These properties differ for Master and Replica. See the Administration Guide for more information.
- Spool properties -- Set the size, minimum disk free size and directory path. See [Spool Directory Settings](#) (see page 43) for more information.
- Event notification properties -- specify a script to run or choose email notification
- Report properties -- choose synchronization or replication reports, specify distribution or script execution
- (Replica) Recovery properties -- set delay or data rewind properties

Settings on the HA Properties Tab

These settings control how switchover and switchback are performed

- Switchover properties -- choose automatic or manual switchover, and reverse replication settings
- Hosts properties -- specify the Master and Replica Fully Qualified Name
- Network Traffic Redirection properties -- choose Move IP, Redirect DNS, or User-defined scripts
- Is Alive properties -- set the heartbeat frequency and check method
- DB Management properties -- instructs Arcserve RHA to manage shares or services on a database server
- Action upon Success properties -- defines custom scripts and arguments for use


If you added a new data or redo log file to the Oracle database, you need to update the scenario. Stop the scenario and rerun Oracle instance auto-discovery by double-clicking the instance name from the Root Directories tab.



Run a Scenario

You can run a single scenario using the following procedure:

To run the scenario

1. From the Scenario pane, select the scenario you want to run.
2. Click Run  on the Standard toolbar.

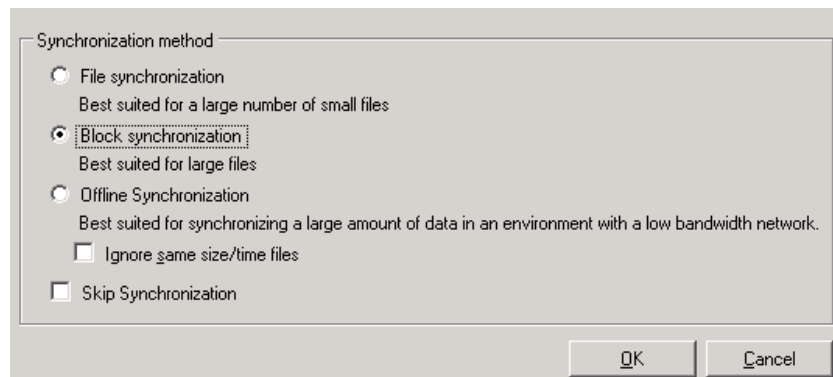
Before initiating synchronization and replication, Arcserve RHA verifies your scenario configuration. When verification completes successfully, the Manager displays the message: *Are you sure you want to run scenario "scenario_name?"* If problems are discovered, the top pane displays any warning and error messages resulting from verification.

Note: Scenario Verification checks many different parameters between the Master and Replica servers to ensure a successful switchover. If any errors or warnings are reported you should not continue until they are resolved.

3. Correct errors before you continue. Errors are reported on the Event pane.

Note: Replication of mount points succeeds only if those were added to the Master before the Engine was started. If you included the mount points in the Master root directories when the Engine was already running, no error is reported but the replication does not start. In this case, you need to restart the Engine on the Master before initiating replication.

When no error is reported, the Run dialog appears and contains synchronization options.




Note: Do not use **Skip Synchronization** for any scenarios replicating a database.

4. If you have a large number of small files, select File Synchronization. If you have large files, select Block Synchronization. If you have low bandwidth, select Offline Synchronization to transfer data to an external device, then perform synchronization from that device. Select the Ignore same size/time files to skip the comparison of files with the same path, name, size and modification time, which are generally identical, to reduce synchronization time. You should enable the Skip Synchronization option only when you are certain the files on both Master and Replica are identical. (The default selections are File Synchronization and Ignore same size/time files option enabled).
5. Click the OK button. Synchronization may take a while, depending on database size and network bandwidth between the Master and Replica. You will receive the following message in the event window when the synchronization is complete: *All modifications during synchronization are replicated.*

At this point, the scenario is operational and active. By default, a Synchronization Report is generated when synchronization finishes. To view the report, refer to the topic, View a Report. You can also generate regular Replication Reports to monitor the replication process on each participating server. For more information, see the *Arcserve RHA Administration Guide*.

Stop a Scenario

To stop a scenario

1. From the Scenario pane, select the scenario you want to stop.
2. To stop the scenario, click the Stop  button on the Standard toolbar.
A confirmation message appears prompting you to approve the scenario stopping.
3. Click Yes in the confirmation message. The scenario stops.
After stopping the scenario, the Manager no longer shows the green play symbol to the left of the scenario, the scenario state turns to Stopped by user and the Statistics tab is no longer available on the Framework pane.

View a Report

Arcserve RHA can generate reports on the replication and synchronization processes. These reports can be stored on your desired location, opened for view from the Report Center, sent by email to a specified address, or they can trigger script execution.

The default storage directory of the generated reports is:
[ProgramFilesFolder]\CA\ArcserveRHA\Manager\reports

To view reports

Note: Though an Exchange report is shown for illustrative purposes, the steps and screens are similar regardless of the type of scenario.

1. To view reports, locate the Tools menu, click Reports, and then select Show Scenario Reports.

The Report Center opens in a new window.

Updated: Tuesday, December 22, 2009 6:39:30 AM

Available Reports per Scenario							
Scenarios							
Scenario Name	Synchronization	Difference	Replication	Assessment Mode	Assured Recovery	Total Reports	
Exchange 1	1	0	0	0	0	1	

Reports								
Drag a column header here to group by that column								
Host	Changes	Date	Time	Type	Summary	Detailed	Size (bytes)	

The Report Center consists of two tables:

- The upper table - Available Reports per Scenario - contains a list of all scenarios that have reports, with the type and number of available reports for each scenario.
 - The lower table - Reports - contains a list of all the reports that are available for the scenario selected in the upper table.
2. To view a specific report, select from the Available Reports per Scenario table the scenario that this report represents. Then, from the Reports table below, click the report you want to open.

Drag a column header here to group by that column								
Host	Changes	Date	Time	Type	Summary	Detailed	Size (bytes)	
XOR\MEXCH267-1	Unknown	Today	03:29:37	Assured Recovery			811	
XOR\MEXCH267-1	Changes found	12/07/06	22:28:48	Synchronization			28415	

Note: Depending on your settings, for Synchronization and Replication reports a Detailed report can be generated in addition to the Summary report. Both reports represent the same process, but the Detailed report also provides a list of the files that participated in the process.

The report you selected appears in a new window.

SYNCHRONIZATION REPORT	
Synchronization mode	BlockSynchronization (include files with the same size and modification time)
Scenario	Exchange 1
Master host	192.168.50.2(1)
Replica host	192.168.50.12(2)
Scenario start time	12/22/2009 06:37:52
Report start time	12/22/2009 06:38:07
Report finish time	12/22/2009 06:39:15

EVENT	BYTES	TIME STAMP	FILE NAME
-------	-------	------------	-----------

Chapter 4: Switching Over and Switching Back

Switchover and *Switchback* is the process in which active and passive roles are exchanged between the Master and Replica servers, so that if the Master is currently active, it changes to passive after Switchover passes the active role to the Replica. If the Replica is active, it changes to passive after Switchover passes the active role to the Master. Switchover can be triggered at the push of a button or automatically by Arcserve RHA when it detects that the Master is unavailable, if you enabled the Perform Switchover Automatically option from the Switchover and Reverse Initiation dialog. When this option is Off, the system notifies you that the Master server is down so you can manually initiate switchover from the Arcserve RHA Manager.

This section contains the following topics:

[How Switchover and Switchback Work](#) (see page 30)

[Initiate Switchover](#) (see page 31)

[Initiate Switchback](#) (see page 33)

[Switchover Considerations](#) (see page 36)

How Switchover and Switchback Work

After the HA scenario starts running and the synchronization process is completed, the Replica checks the Master on a regular basis, by default every 30 seconds, to see if it is alive. There are three types of monitoring checks:

- **Ping** -- a request sent to the Master to verify that the Master is up and responding
- **Database check** -- a request that verifies the Oracle Server is running
- **User-defined check** -- a custom request you can tailor to monitor specific applications

If an error occurs with any part of the set, the entire check is considered to have failed. If all checks fail throughout a configured timeout period (by default, 5 minutes), the Master server is considered to be down. Then, depending on the HA scenario configuration, Arcserve HA sends you an alert or automatically initiates a switchover.

When you created an HA scenario, you defined how you want the switchover to be initiated.

- If you selected the Initiate Switchover manually option from the Switchover and Reverse Replication Initiation page, perform a manual switchover. For more information, refer to the topic, [Initiate Switchover](#) (see page 31).
- If you selected the Initiate Switchover automatically option, you can still perform a manual switchover, even if the Master is alive. You can initiate switchover when you want to test your system, or you want to use the Replica server to continue the application service while some form of maintenance is performed on the Master server.

Triggered (automatic) switchover is in all ways identical to manual switchover performed by the administrator. Automatic switchover is triggered by a resource failure on the master server rather than by an administrator manually initiating the switchover by clicking the Perform Switchover button. Server ping response, application service status, and database connectivity are monitored. The timeout parameters are configurable and are more extensively covered in the *Arcserve HA Administration Guide*

When you created an HA scenario, you defined how you want the reverse scenario to be initiated.

- If you selected the Initiate Reverse Replication automatically option from the Switchover and Reverse Replication Initiation page, replication in the reverse direction (from Replica to Master) automatically begins after a switchover, if the original Master server is online.
- If you selected the Initiate Reverse Replication manually option, you need to perform switchback manually. If you performed switchover with the Master server offline, or if you stop a backward scenario and restart it again, in both cases, data from the Replica to the Master must be resynchronized before you can switch back to the Master.

When the Reverse Replication feature is off, to start reverse replication after a switchover has occurred, click the Run button. The benefit to this feature is, if both the master and replica servers were online and connected during switchover, resynchronization in the reverse direction is not required. Resynchronization involves comparing the data on the master and replica servers to determine which changes to transfer before real-time replication starts; this can take some time. If automatic reverse replication is turned on, and both servers were online during switchover, replication is reversed without the need for resynchronization. This is the one situation in which resynchronization is not required.

Initiate Switchover

Once triggered, whether manually or automatically, the switchover process itself is fully automated.

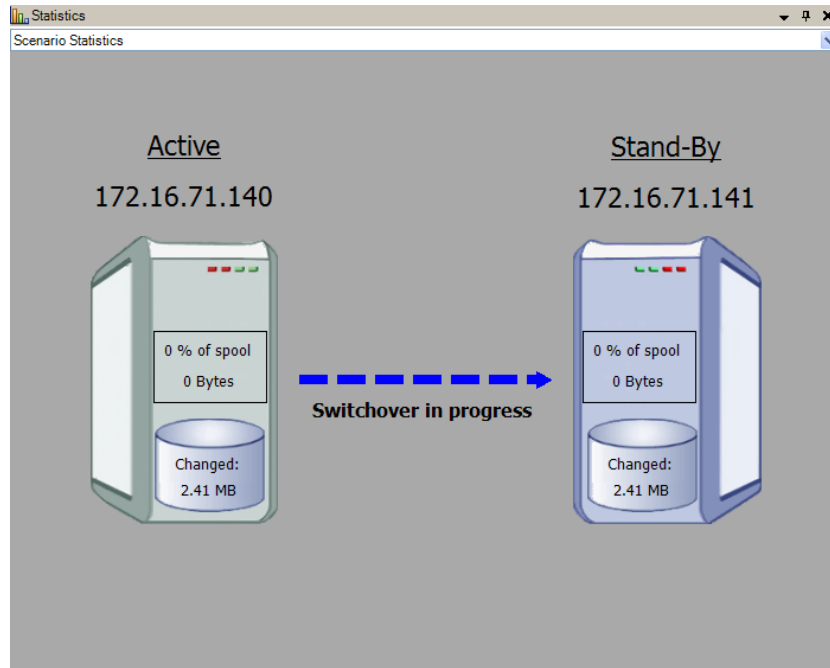
Note: Though the following steps show Exchange scenario screens as examples, the procedure is similar for all server types.

To initiate manual switchover

1. Open the Manager and select the desired scenario from the Scenario pane. Ensure that it is running.
2. Click Perform Switchover.
A confirmation message appears.

3. Click OK.

A switchover from the Master server to the Replica server is initiated.



Detailed information about the switchover processes is displayed in the Events pane during switchover.

After the switchover is completed the scenario stops.

HA Scenarios				
Scenario	State	Product	Server	Mode
MS Exchange HA	Stopped on Aut...	HA	Exchange	Online
Hosts				
172.16.71.140	Changed	Synchronized	Files	In pool
172.16.71.141				

Note: The only case in which the scenario may continue to run after switchover is when you have selected Start automatically for Automatic Reverse Replication.

In the Event pane a message appears, informing you that switchover completed and the scenario has stopped.

Now, the Master becomes the stand-by server and the Replica becomes active server.

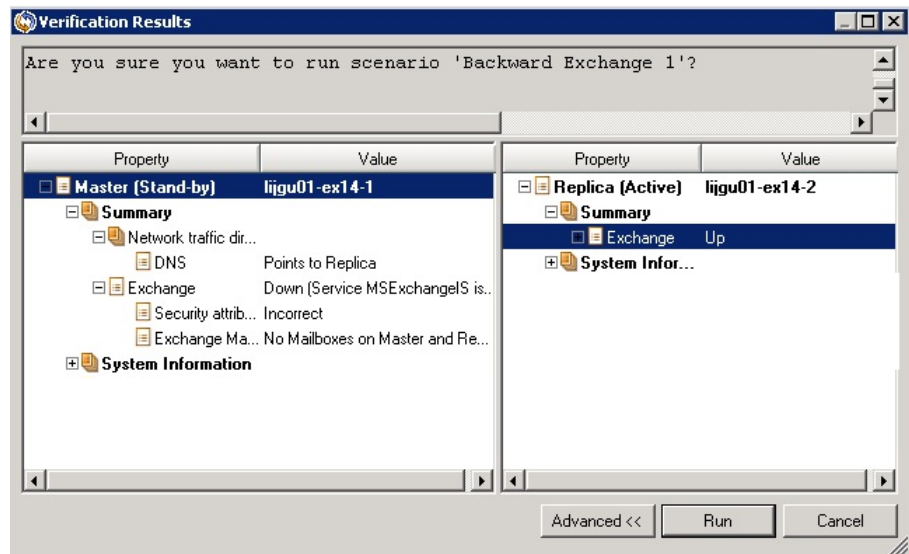
Initiate Switchback

After a switchover is initiated, whether manually or automatically, at some point, you will want to reverse the server roles and make the original Master the active server again the Replica the standby server. Before you switch back the roles between servers, decide if you want to the data on the original Replica server to overwrite the data on the original Master. If yes, you must first perform a reverse scenario, called a backward scenario

To initiate switchback

1. Ensure that both Master and Replica servers are available on the network and that the Arcserve RHA Engine is running.
2. Open the Manager and select the desired scenario from the Scenario pane.
3. Perform one of the following:
 - If the scenario is already running, skip directly to Step 4
 - If the scenario is not running, perform these steps and then go to Step 4:
 - a. Click Run on the toolbar to start the scenario.

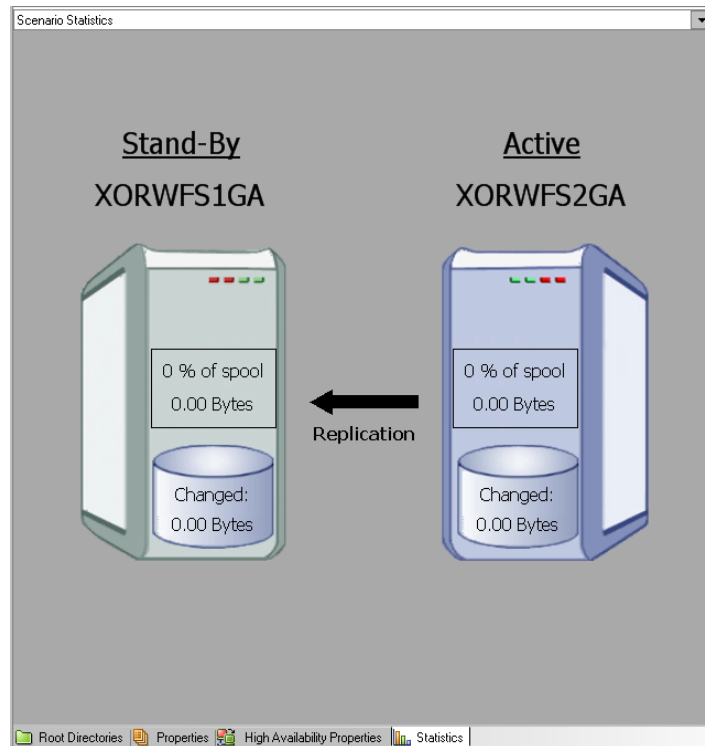
Arcserve HA detects that a switchover has occurred and verifies its state and configuration. After verification completes, the Verification Results dialog appears, listing existing errors and warnings if detected, and prompting you to approve the running of the backward scenario. If desired, click the Advanced button to open an additional pane with detailed information about the hosts that participate in the scenario.



- b. Select a synchronization method from the Run dialog and click OK to start resynchronization.

Note: For Oracle servers, choose Block Synchronization.

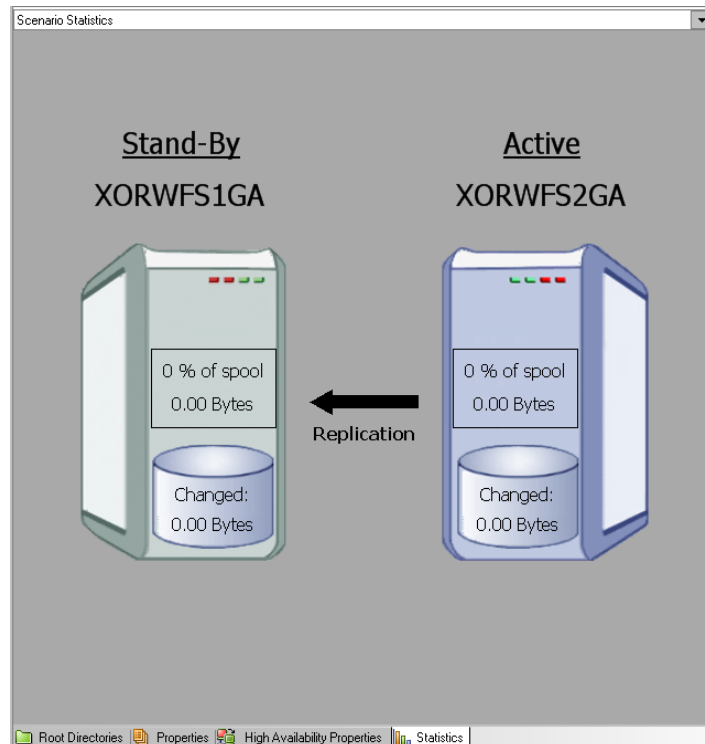
After resynchronization completes, you receive a message in the Event pane: All modifications during synchronization period are replicated. Now, replication from the active server to the standby server begins:



Note: You are now ready to reverse the roles between the Master and Replica servers.

4. Click Perform Switchover on the toolbar while the scenario is running to reverse the server roles. A confirmation message appears.
5. Click Yes to clear the message and start the switchback process.

After the switchback is completed, the server roles are reversed back and the scenario automatically stops.



Note: The scenario will continue to run after the switchback when the Reverse Replication Initiation option is defined as Start Automatically.

You may now run the scenario again in its original (forward) state.

Switchover Considerations

To prevent overwriting data, the best practice is to set *either* the Switchover or the Reverse Replication Initiation property to Automatic. If a server fails while both properties are set to Automatic, Arcserve RHA triggers Switchover without administrative involvement and could start Reverse Replication before you have investigated the cause of the failure. During Reverse Replication, Arcserve RHA overwrites data on your production server.

If a crash or outage occurs during switchover, you may need to perform the Recover Active Server procedure.

Chapter 5: Recovering Data

This section contains the following topics:

[The Data Recovery Process](#) (see page 37)

[How to Recover Oracle Data](#) (see page 37)

[Setting Bookmarks](#) (see page 38)

[Data Rewind](#) (see page 39)

[Recover Active Server](#) (see page 41)

The Data Recovery Process

When an event causes loss of Master data, the data can be restored from any Replica. The recovery process is a synchronization process in the reverse direction - from a Replica to the Master.

Arcserve RHA enables you to recover data in two ways:

- **Recover lost data from the Replica to the Master** -- this option is a synchronization process in the reverse direction and requires you to stop the scenario. (This option is not recommended for Oracle, SQL or Exchange scenarios.)
- **Recover lost data from a certain event or point in time (Data Rewind)** -- This option uses a process of stamped checkpoints and user-defined bookmarks to roll corrupt data on the Master back to a time before corruption occurred.

Important! You must stop replication to initiate recovery.

How to Recover Oracle Data

Data recovery is covered in full detail in the Arcserve RHA Administration Guide.

Because the **Recover Lost Data from Replica** procedure requires that you stop Oracle instances on both Master and Replica servers, it is not recommended and therefore, not included in this Operation Guide.

Instead, we recommend that you switch to the Replica, and run a backward scenario to synchronize data from the Replica back to the Master. In this manner, stopping the Oracle instance on the active Replica server is not required.

If you find that data on the Master is corrupted, you can use the Data Rewind process to restore it, if desired. That procedure is provided in this Guide.

Setting Bookmarks

A *bookmark* is a checkpoint that is manually set to mark a state back to which you can revert. We recommend setting a bookmark just before any activity that can cause data to become unstable. Bookmarks are set in real-time, and not for past events.


Notes:

- You can use this option only if you set the Recovery--Data Rewind option to *On* (default setting is *Off*).
- You cannot set bookmarks during the synchronization process.
- You can insert manual bookmarks for Full System HA scenarios.

To set a bookmark

1. Select the Replica host on the Scenario pane from which you want to rewind data when the required scenario is running.
2. Select the Set Rewind Bookmark option on the Tools menu.

The Rewind Bookmark dialog opens.



The text that appears in the Rewind Bookmark dialog appears in the Rewind Points Selection dialog as the bookmark's name. The default name includes date and time.

3. Accept the default name, or enter a new name for the bookmark, and click OK.

Note: It is recommended that you provide a meaningful name that will later help you recognize the required bookmark.

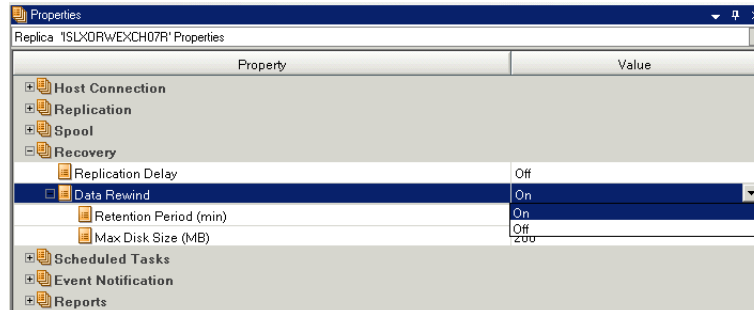
The bookmark is set.

Note: In some scenarios, such as Full System HA, applying journal changes is suspended until the bookmark is created and then resumed.

Data Rewind

The Data Rewind recovery method allows you to rewind data to a point in time before it was corrupted. The rewind process takes place on the Replica server before the reverse synchronization process starts. The Data Rewind method uses rewind points or bookmarks that enable you to reset the current data back to a previous state.

You can use this option only if you set the **Recovery - Data Rewind** option to **On**.



If this option is set to Off, the system will not register data rewind points. For more information about Data Rewind parameters (Retention Period, Max Disk Size), see the *Arcserve RHA Administration Guide*.


Important! The data rewind process operates in one way only - there is no replay forward. After rewind, all data subsequent to the rewind point will be lost, since data after the rewind point will be overwritten with new data.


Note: The automatic registration of the rewind points starts only after the synchronization process is completed, and the message **All modifications during synchronization period are replicated** appears on the Event pane. Similarly, you cannot manually set bookmarks during synchronization.

To recover lost data using rewind points

1. On the Manager, from the Scenario pane select the desired scenario and stop it.
2. [For database applications only] stop the database services on the Master host.
3. On the Manager, from the scenario folder select the Replica host:

Note: If multiple Replica servers participate in the required scenario, select the Replica from which you want to recover data.



- From the **Tools** menu, select **Restore Data**, or click the **Restore Data**  button. If you are prompted for user credentials, enter the appropriate information and click OK.

The **Recovery Method** page of the Restore Data Wizard opens.

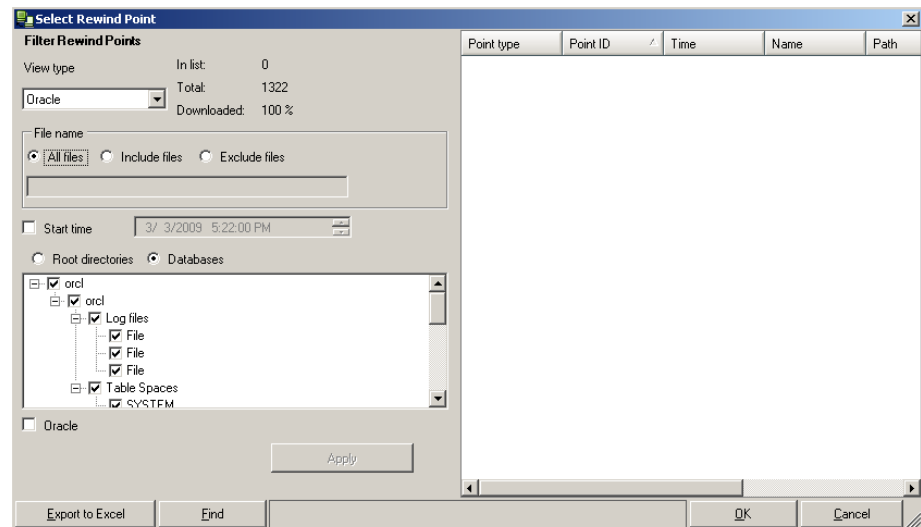
- Select one of the Rewind data options, depending on whether you want the rewind data synchronized back to the Master (option 2) or left on the Replica only (option 3).

Note: If the user credentials you used to log in to the Manager are different than the ones required for working with the Engine on the Replica, a **User credentials** dialog appears, asking you to enter log on account details for the selected Replica.

After you select a Rewind data option, a Recovery scenario is automatically created. This Recovery scenario will run until the end of the rewind process.

- Click **Next**. The **Rewind Point Selection** page is displayed.
- Wait until the **Select Rewind Point** button is enabled, and click it to view the existing rewind points.

The **Select Rewind Point** dialog opens.



The **Select Rewind Point** dialog displays a list of all rewind points. These include modifications of folders and files that were automatically registered by the system and user-defined bookmarks.

The list can be filtered according to the rewind point type or other criteria, using the **Filter Rewind Points** pane on the left.

- Select the required rewind point, and click **OK**.

Note: If you want to use a Bookmark as the rewind point, it is best practice to select the closest rewind point that indicates an actual event.

You return to the **Rewind Point Selection** page, now displaying information about the rewind point you selected.

9. Click **Next**. The **Synchronization Method** page is displayed.
10. Select the **Block Synchronization** method and click **Finish**.

Note: If the user credentials you used to log in to the Manager are different than the ones required for working with the Engine on the Replica, a **User credentials** dialog appears, asking you to enter log on account details for the selected Replica.

Arcserve RHA rewinds the data to the point you selected. After the rewind process ends, you receive the following message in the Event pane: **Rewind process is completed successfully**.

If you chose to replace the data on the Master with the data on the Replica, Arcserve RHA starts a synchronization process from the Replica to the Master. Once the process ends, the temporary Recovery scenario is stopped and then deleted.

11. By default, once a data recovery occurs a Synchronization Report is generated.
Now, the Replication process can restart on the original scenario.

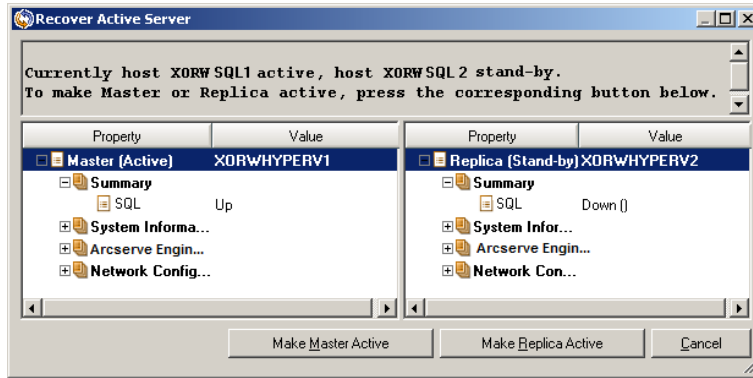
Recover Active Server

In certain circumstances, it may be necessary to forcibly make the Master or Replica server the active server without completing the data synchronization process. For example, if switchover occurred but no data was changed on the Replica server. In this case you may even have newer data on the Master server making it undesirable to synchronize data from the Replica to the Master server. Arcserve HA allows for this option through a process called Recover Active Server. To use this option, ensure that the scenario is stopped, and select *Recover Active Server* from the Tools menu.

Important! While this option is the right choice in many situations, use it with caution. If used improperly data loss can occur. Normally, Arcserve HA will not allow switchover from one host to another until all data is synchronized. It is designed this way so users are not redirected to an out of date data set that then overwrites what may be a more current data set. When using Recover Active Server, Arcserve HA is forcing users to one server or the other with no regard as to which server has the correct data set. Thus, as an administrator, you must manually ensure that the server you are making active has the most up to date data set.

If Recover Active Server does not solve the problem, you can manually recover a server. For more information, refer to the topic, [Manually Recover a Failed Server](#). (see page 44)

Select either *Make Master Active* or *Make Replica Active* depending onto which server you want to force the active role.



Important! If a legitimate switchover in a disaster situation occurs and users are redirected to the Replica server for any period of time, it is important to replicate all changes on the Replica back to the Master before making the Master server active. Using *Recover Active Server* in such a situation results in loss of data.

Appendix A: Additional Information and Tips

This section contains the following topics:

[Spool Directory Settings](#) (see page 43)

[Manually Recover a Failed Server-Move IP Address](#) (see page 44)

[Troubleshooting Oracle Databases](#) (see page 44)

[Cannot Connect Oracle Enterprise Manager After Switchover](#) (see page 45)

Spool Directory Settings

The Arcserve RHA spool is a folder on disk where data to be replicated is backed up (spooled) if bandwidth is not sufficient to transfer the amount of changes in real-time. Data can spool due to temporary network disconnections, network congestion, or simply because the network bandwidth is not sufficient to transfer the amount of data changing over on the server. In addition to storing changes waiting on available bandwidth, spool space is also used as part of the normal synchronization process. Thus, some spool build up during synchronization is normal.

Place the spool folder on a drive with relatively low use such as a dedicated volume or boot/system volume. Do not place the spool folder on a volume containing frequently accessed system (OS), user, or application data. Examples include volumes containing databases, shared files, or the system pagefile. By default, the spool folder is located in the tmp folder under the Arcserve RHA installation directory. The spool parameters, located in the properties tab (on both master and replica) or set with the New Scenario Wizard, determines how much disk space is available for the spool. In most cases the default values are sufficient. However, if you change this value, it should be at least 10% of the total dataset size. For example, if you are replicating 50 GB of data on a server you should ensure that at least 5 GB of space is available for spool.

Important! If you change the spool location, remember to remove the new path from file level antivirus scans: both scheduled and real time.

Note: The Spool Directory is not a pre-allocated space folder and will be used only if needed.

Manually Recover a Failed Server-Move IP Address

If IP Redirection is used, you must remove the IP address manually. You cannot use this method for scenarios that do not support Move IP redirection (Hyper-V HA, CS HA).

To recover a failed server using Move IP Address redirection method

1. Boot the Master server without a network connection, to avoid IP conflicts.
2. From the TCP/IP properties dialog, remove the additional IP address.
3. Reboot the server and reconnect to the network.
4. If it is not already running, start the scenario from the Manager. If automatic reverse replication was set to On, the scenario runs in backward mode so that the Replica server is now active and the Master server is on standby.
5. Wait for synchronization to complete.
6. Perform a manual switchover to return the active role to the Master server. It is recommended that you do so outside of normal business hours.

Troubleshooting Oracle Databases

Oracle database fails to start after switchover

Symptom:

My vCenter Server HA scenario uses an Oracle database. After switchover, the Oracle database does not restart and I get the following errors:

- ER00319, 83, Error, <HOSTNAME><DATE/TIME>, Switchover Aborted
- ER00303, 82, Error, <HOSTNAME><DATE/TIME>, Unable to start vCenter Services
- ER00360, 81, Error, <HOSTNAME><DATE/TIME>, Unable to start vCenter after switchover. Service 'vpxd' not started

Solution:

These errors occur when the Oracle database does not successfully mount the following switchover. Use the command line to solve the problem:

1. Shut down

```
[ORACLE_HOME]\bin\oradim.exe -shutdown -sid orcl -usrpwd * -shutmode immediate
```
2. Start again

```
[ORACLE_HOME]\bin\oradim.exe -startup -sid orcl -usrpwd * -nocheck 0
```

Cannot Connect Oracle Enterprise Manager After Switchover

Symptom

After a switchover or failover occurs, the Oracle Enterprise Manager on the Replica is not accessible.

Solution

Reconfigure the Database Console. Before you begin, stop the running scenario or suspend the Is-Alive check. The following procedure is an example. Provide the correct values for <oracle_db_instance_name>, Oracle Database Version (the example uses 11g R2), file paths, and SID for your environment.

1. set oracle_home=D:\app\administrator\product\11.2.0\dbhome_1
2. set oracle_sid=orcl
3. emca -deconfig dbcontrol db -repos drop
4. To verify the TEMP DB already exists, perform the following steps:
 - a. sqlplus /nolog
 - b. connect / as sysdba
 - c. select * from dba_temp_files where tablespace_name = 'TEMP';

If an error occurs, perform these steps:

- alter database temp file
'D:\APP\ADMINISTRATOR\ORADATA\ORCL\TEMP01.DBF' drop including
datafiles;
 - alter tablespace temp add tempfile
'D:\APP\ADMINISTRATOR\ORADATA\ORCL\TEMP01.DBF';
5. emca -config dbcontrol db -repos create
 6. net stop OracleDBConsole<oracle_db_instance_name>
 7. net start OracleDBConsole<oracle_db_instance_name>

Index

A

- About This Guide • 7
- Add IP on Master Server • 12
- Additional Information and Tips • 43
- Arcserve Product References • 3

C

- Cannot Connect Oracle Enterprise Manager After Switchover • 45
- Considerations • 14
- Contact Arcserve • 3
- Create Oracle HA and DR Scenario • 15
- Creating and Using Scenarios • 13

D

- Data Rewind • 39
- DNS Redirection • 11
- Documentation Changes • 4

H

- How Redirection Works • 11
- How Switchover and Switchback Work • 30
- How to Recover Oracle Data • 37

I

- Initiate Switchback • 33
- Initiate Switchover • 30, 31
- Introduction • 7

L

- Log On Account Conditions • 8

M

- Manually Recover a Failed Server-Move IP Address • 41, 44
- Move IP Redirection • 12

R

- Recover Active Server • 41
- Recovering Data • 37
- Redirection Methods • 11
- Related Documentation • 7

- Run a Scenario • 23

S

- Scenario Properties • 21
- Scripts Redirection • 12
- Server Requirements • 8
- Servers Operating in a Workgroup • 9
- Setting Bookmarks • 38
- Spool Directory Settings • 22, 43
- Stop a Scenario • 25
- Switching Over and Switching Back • 29
- Switchover Considerations • 36

T

- The Data Recovery Process • 37
- Troubleshooting Oracle Databases • 44

V

- View a Report • 25