

Arcserve® Replication and High Availability

Microsoft SQL Server Operation Guide 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
Chapter 2: Protecting Microsoft SQL Server Environments	9
Server Requirements	9
Base Configuration	9
SQL Server Configuration Requirements	10
Log On Account Conditions	10
SQL Servers Operating in a Workgroup	11
About Clusters	11
Register Arcserve RHA Licenses	11
Chapter 3: Creating Replication and High Availability Scenarios	13
Create SQL Server Replication Scenarios	13
Create SQL Server HA Scenarios	14
Manage Services	16
Chapter 4: Redirection Methods	19
How Redirection Works	19
DNS Redirection	19
Move IP Redirection	20
Add IP on the Master Server	20
Cluster Move IP	24
Switch Computer Name Redirection	27
Automatic Redirection Using Switch Computer Name	28
Scripts Redirection	28
Chapter 5: Managing Replication and High Availability Scenarios	29
Scenario Properties	29
Run the Scenario from Outside the Wizard	32
View a Report	33
Stop a Scenario	34

Chapter 6: Switching Over and Switching Back	35
How Switchover and Switchback Work.....	35
Initiate Switchover	37
Initiate Switchback	38
Switchover Considerations.....	41
Chapter 7: Recovering Data	43
The Data Recovery Process	43
Recover Lost Data from Replica	43
Setting Bookmarks	45
Data Rewind	46
Recover Active Server	48
Appendix A: Additional Information and Tips	49
Spool Settings.....	49
Rename the Microsoft SQL Server 2005	50
Recovering Servers.....	50
Manually Recover a Failed Server-Move IP Address.....	51
Manually Recover a Failed Server-Switch Computer Name	52
Manually Recover Failed Server-IP and Switch Computer Name	53
Index	55

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 Microsoft SQL Server 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)

About This Guide

This document describes how you can implement a Arcserve RHA solution for Microsoft SQL Server. 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*.

Chapter 2: Protecting Microsoft SQL Server Environments

This section contains the following topics:

[Server Requirements](#) (see page 9)

Server Requirements

To implement high availability for SQL servers, refer to the following list of requirements. If you do not have the license required to access support for a specific server type, please contact Tech Support.

Base Configuration

Base 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.

- All IP addresses are statically assigned (DHCP-assigned IP addresses on the Master or Replica server are not supported)
- The protected server is not a domain controller or DNS server.
- (In the Active Directory environment) Both Master and Replica servers should reside in the same Active Directory forest and also be members of the same domain or trusted domains.

SQL Server Configuration Requirements

One or more instances of supported Microsoft SQL Server is installed on each server. For supported SQL servers, see the Arcserve RHA Release Notes.

- Both servers must have the same SQL version, service packs and hot fixes installed
- Both servers must hold identical SQL server instances (default or named)
- (In the Active Directory environment) Both Master and Replica servers must reside in the same Active Directory forest and be members of the same domain or trusted domains
- Drive letters containing database files must be identical on both servers
- The full path to the default system database of each instance must be identical on both servers
- Verify that the port defined in the Network Configuration TCP/IP properties of the SQL instances is assigned statically and is identical on both Master and Replica
- When the SQL account is Network, make sure that you create a database login for the NT Authority\Network Service account. See Microsoft documents for more details.

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.
Note: In MS SQL server 2012, Local System (NT AUTHORITY\SYSTEM) is not automatically provisioned in the sysadmin server role. See Microsoft documents on how to provision the sysadmin server role to the NT AUTHORITY\SYSTEM account. Alternatively, use Administrator account to install and log in to the engine service.
- When your SQL servers are in a workgroup, enable "sa" account on Master and Replica server before you run the scenario.

SQL Servers Operating in a Workgroup

For servers in a workgroup, set the Arcserve HA Engine service account to a user that is a member of the Local Administrators group. Servers in a workgroup can use Redirect DNS only with DNS servers that allow non-secure updates. You can use Move IP, Switch Computer Name, and custom redirection scripts normally.

When Switch Computer Name is used, you may use the System Account or the Local Administrators group as long as the account was added to the Microsoft SQL logins.

About Clusters

Installing on clusters is much the same as a standard installation. To configure Arcserve RHA on a cluster, enter the Virtual Server Network Name (or IP Address) resource (in the group you intend to protect) as the Master or Replica name. Do not use node names or IP addresses when configuring the scenario. Also, you must install the Arcserve RHA Engine to all cluster nodes (see *Server Setup*).

The only configuration that requires some preparation is the use of IP Move in conjunction with a cluster. For detailed instructions on how to use Move IP with clusters, please see [Cluster Move IP](#) (see page 24).

Register Arcserve RHA Licenses

The Arcserve RHA licensing policy is based on a combination of several parameters which include the following:

- the operating systems involved
- the required solution
- the supported application and database servers
- the number of participating hosts
- additional modules (for example, Assured Recovery)

The license key that is generated for you is therefore tailored to your exact needs.

After logging in for the first time, or if your old license has expired, you must register the Arcserve RHA product using your license key. To register the product, you need to open the Manager, which does not depend on the existence of a valid registration key. After the Manager opens, a License Warning message appears, prompting you to register the product. A License Warning message also appears when your license is about to expire during the next 14 days.

When you are creating a scenario, some of the options might be disabled following the terms of your license. However, you can create any number of scenarios, since the validity of your license key is confirmed, before you try to run a specific scenario. Only when you click the Run button, the system checks whether you are allowed to run the selected scenario according to your license key. If the system determines that you do not have the required license for running this scenario, the scenario will not run and a message will appear on the Event pane informing you of the type of license you need.

To register Arcserve RHA using the license key

1. Open the Manager.

The Welcome message appears, followed by a License Warning message informing you that your product is not registered. You are prompted to register it.

2. Click OK to close the message.
3. Open the Help menu and select the Register option.

The Register Arcserve RHA dialog opens.

4. Complete the following fields:

- Registration Key field - enter your registration key.
- [Optional] In the Company Name field - enter your company name

5. Click the Register button to register the product and close the dialog.

You can now start working with the Arcserve RHA Manager according to your license permissions.

Chapter 3: Creating Replication and High Availability Scenarios

This section describes the most common setup for Arcserve HA for Microsoft SQL Server.

This section contains the following topics:

[Create SQL Server Replication Scenarios](#) (see page 13)

[Create SQL Server HA Scenarios](#) (see page 14)

[Manage Services](#) (see page 16)

Create SQL Server Replication Scenarios

Creating scenarios is covered in full detail in the *Arcserve RHA Administration Guide*. This section provides additional information specific to a Microsoft SQL Server disaster recovery scenario. The Scenario Creation Wizard guides you through the steps required to create a disaster recovery scenario. When completed, you should run your scenario to start data synchronization. Synchronization could take a while, depending on database size and network bandwidth. Once synchronization completes, your scenario now maintains the Replica server so that it can take over for the Master the moment a failure is detected.

You can also protect MS SQL BLOB data. When the master server has MS SQL server enabled FILESTREAMS, the RHA engine discovers and replicates all BLOB data. The engine discovers and replicates even if the BLOB data is stored in locations other than the database files. The root directory also adjusts automatically according to the BLOB data path.

Please read the entire procedure, including cross-referenced information, if applicable, before you proceed.

To create a SQL scenario

1. From the Arcserve RHA Manager, choose Scenario, New or click the New Scenario button.
2. When the Welcome dialog opens, select Create New Scenario and click Next.

3. When the Select Scenario Type dialog opens, select SQL, Disaster Recovery Scenario, and Replica Integrity Testing for Assured Recovery (optional). For more information on Assured Recovery, see the *Arcserve RHA Administration Guide*.
4. When the Master and Replica hosts dialog opens, name your scenario and provide the hostname or IP address for the Master and Replica servers. If either server is a MSCS cluster, enter the cluster resource virtual server name or IP address. Click Next. For more information, see [Redirection Methods](#) (see page 19).
5. Wait for Engine Verification to complete and click Next. If needed, click Install to upgrade the Engine on one or both servers and then click Next.

The Database for Replication dialog opens, listing all auto-discovered results for the specified Master. By default, all databases are included.
6. Change selections, as desired and click Next.
7. When the Scenario Properties dialog opens, configure additional properties, if needed. If you use NTFS ACLs with domain accounts for user access control, we recommend that you choose the Replicate NTFS ACL option and click Next. For more information, see Scenario Properties or the *Arcserve RHA Administration Guide*.

The Master and Replica Properties dialog opens.
8. Accept default settings or make the desired changes and click Next.
9. Click Next to initiate scenario verification. If errors are reported, you should resolve them before continuing. At successful verification, click Next to complete scenario creation.
10. Choose Run Now or Finish, as desired. Run Now starts synchronization. Finish allows you to run the scenario later. See [Run the Scenario from Outside the Wizard](#). (see page 34)

Create SQL Server HA Scenarios

Creating scenarios is covered in full detail in the *Arcserve RHA Administration Guide*. This section provides additional information specific to a Microsoft SQL High Availability scenario. The Scenario Creation Wizard guides you through the steps required to create a high availability scenario. When completed, you should run your scenario to start data synchronization. Synchronization could take a while, depending on database size and network bandwidth. Once synchronization completes, your high availability scenario now maintains the Replica server so that it can take over for the Master the moment a failure is detected.

Please read the entire procedure, including cross-referenced information, if applicable, before you proceed.

To create a SQL HA scenario

1. From the Arcserve RHA Manager, choose Scenario, New or click the New Scenario button.
2. When the Welcome dialog opens, select Create New Scenario and click Next.
3. When the Select Scenario Type dialog opens, select SQL, High Availability Scenario, and Replica Integrity Testing for Assured Recovery (optional). For more information on Assured Recovery, see the *Arcserve RHA Administration Guide*.
4. When the Master and Replica hosts dialog opens, name your scenario and provide the hostname or IP address for the Master and Replica servers. If either server is a MSCS cluster, enter the cluster resource virtual server name or IP address. Click Next. For more information, see [Redirection Methods](#) (see page 19).
5. Wait for Engine Verification to complete and click Next. If needed, click Install to upgrade the Engine on one or both servers and then click Next.

The Database for Replication dialog opens, listing all auto-discovered results for the specified Master. By default, all databases are included.

Note: You can only de-select user created databases.

6. Change selections, as desired and click Next.
7. When the Scenario Properties dialog opens, configure additional properties, if needed. If you use NTFS ACLs with domain accounts for user access control, we recommend that you choose the Replicate NTFS ACL option and click Next. For more information, see Scenario Properties or the *Arcserve RHA Administration Guide*.

The Master and Replica Properties dialog opens.
8. Accept default settings or make the desired changes and click Next.
9. Wait for the Switchover Properties dialog to retrieve information. Configure the desired redirection properties and click Next. For more information, see [Switching Over and Switching Back](#) (see page 35).
10. From the Switchover and Reverse Replication Initiation dialog, choose automatic or manual switchover, and automatic or manual reverse replication, as needed.

Setting both of these options to automatic is not recommended. For more information, see Scenario Properties or the *Arcserve RHA Administration Guide*.
11. Click Next to initiate scenario verification. If errors are reported, you should resolve them before continuing. At successful verification, click Next to complete scenario creation.
12. Choose Run Now or Finish, as desired. Run Now starts synchronization. Finish allows you to run the scenario later. See [Run the Scenario from Outside the Wizard](#). (see page 34)

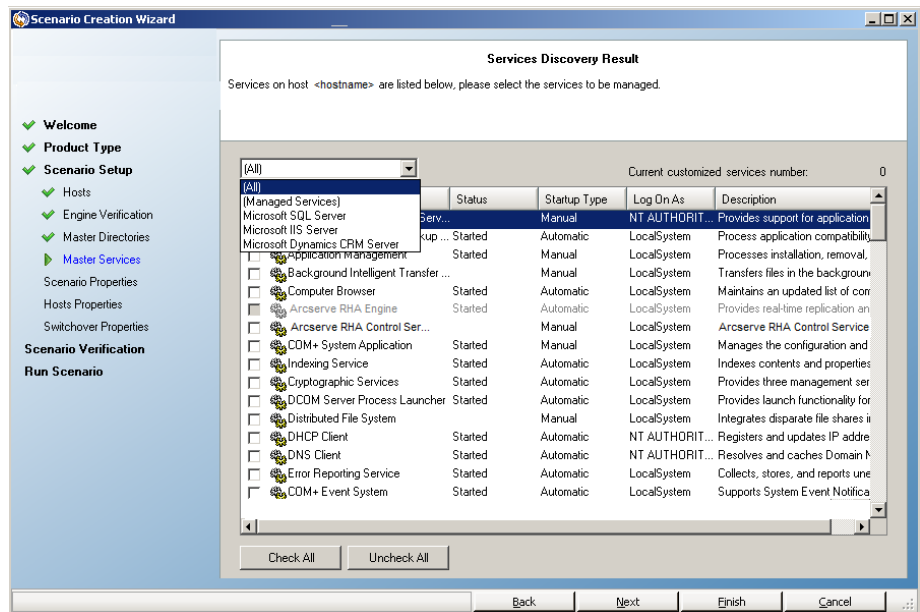
Manage Services

As part of scenario creation or modification, you can specify the services to manage. During scenario creation, the services management screens are displayed in the Scenario Creation Wizard. For existing scenarios, you can also manage services from the Arcserve RHA Manager Root Directories tab.

Services discovered on the specified Master server are automatically shown on the Services Discovery Result screen in the Scenario Creation Wizard.

The following steps are for Custom Application scenarios.

To manage services



- **All** - lists all services discovered on the Master server
- **Managed Services** - lists only the checked services
- **Oracle Database** - lists Oracle-related services if the current host has Oracle installed
- **Microsoft SQL Server** - lists SQL Server-related services if the current host has SQL Server installed
- **Microsoft IIS Server** - lists IIS Server-related services if the current host has IIS Server installed
- **Microsoft SharePoint Server** - lists SharePoint Server-related services if the current host has SharePoint Server installed
- **VMware vCenter Server** - lists vCenter Server-related services if the current host has vCenter Server installed


- **Microsoft Exchange Server** - lists Microsoft Exchange Server-related services if the current host has Microsoft Exchange Server installed
 - **Microsoft Dynamics CRM Server** - lists Microsoft Dynamics CRM Server-related services if the current host has Microsoft Dynamics CRM Server installed
1. Select a Service to Monitor. Click the box to the left of each service listed to select it for monitoring.

Important! Do not use Services Management to monitor every service on the Master server in a single scenario. This scenario type is not designed to protect an entire server.

2. Click Next to proceed to the Services Setting screen.

Services Setting

Managed services are listed below, please set the properties for each service.



3. In the Start Order column for each service you chose, specify the numeric value representing start order. For services where order does not matter, use the default value, (Not Set). The options available in the drop down list update as you configure the value. The first service has only two options: Not Set and 1. The second service has three options: Not Set, 1 and 2, and so on. If you assign the same start order to two services, Arcserve RHA automatically reorders the selections you already made.
4. In Replication scenarios, the Critical column is disabled. In HA scenarios, use the Critical column to specify if a service should trigger switchover when it fails. By default, all services are marked Critical. Clear the box for any service whose failure does not require switchover to the stand-by server.

Chapter 4: Redirection Methods

This section contains the following topics:

[How Redirection Works](#) (see page 19)

[DNS Redirection](#) (see page 19)

[Move IP Redirection](#) (see page 20)

[Switch Computer Name Redirection](#) (see page 27)

[Scripts Redirection](#) (see page 28)

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 Microsoft SQL Server.

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

DNS Redirection changes the DNS "A" Record of the Master server to resolve to IP address of the Replica server. 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 RHA to redirect any DNS A (host) record via the *Master's name in DNS* setting in the switchover properties tab.

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.

Add IP on the Master Server

You need to add an additional IP address to the Master host, (which is denoted as *Arcserve-IP* in the following steps) to use Move IP redirection in your HA scenarios. This new IP address is used for Arcserve RHA internal communication and replication. This is necessary because once switchover occurs, the current production IP address is no longer available on the Master -- it switches to the Replica server.


Important! Perform the following only if you are using the Move IP redirection method.

To add IP Address to Master Server

1. Open the Control Panel and select Network Connections.
2. Right-click Local Area Network and select Properties.
3. Click Internet Protocol (TCP/IP) and then click the Properties button.
4. Click Advanced.

5. Click Add and enter an additional IP address (Arcserve-IP).

In the following screenshot, the Arcserve-IP IP address is 192.168.220.23 and the current production server IP address is 192.168.220.111.



6. Click Add.
7. Click OK.
8. Click OK to exit the LAN settings.

After you add the IP to the Master, you must add the Arcserve-IP to your HA scenarios. There are two ways to add the Arcserve-IP address to an HA scenario:

- For new scenarios, from directly in the Wizard
- For existing scenarios, by modifying the master host name


The procedures for both ways follow.

Add Arcserve-IP to Existing Scenarios-SQL

Perform this procedure only if you are using the Move IP redirection method.

To add the Arcserve-IP to existing scenarios:

1. On the Scenario pane, select the required Master host.



2. Right-click the Master and select **Rename** from the pop-up menu. Then, enter the Arcserve-IP address.
3. On the Framework pane, select the **Switchover** tab and then select the Replica server as the switchover host.

- Set the **Move IP** option to On. Ensure that the IP address under **Move IP, IP/Mask** matches the production server IP address: this is the IP address that will switch over. If you are moving more than one IP address you can add multiple production IP addresses by selecting **Click here to add new IP/Mask**.

The screenshot shows the 'High Availability Properties' dialog box. The 'Move IP' option is set to 'Off'. Below it, 'Redirect DNS' is set to 'On', and 'DNS Servers IPs' is set to 'Off'. The 'DNS Servers IPs' section is expanded, showing a table of DNS IP addresses. The first entry is '141.202.226.10', and the second is 'Click here to add new IP.'. The 'DNS TTL (sec)' is set to '60', and 'Active Directory Integrated' is set to 'On'. The 'Master IPs in DNS' section is also expanded, showing a table of IP addresses. The first entry is '141.202.226.74', and the second is 'Click here to add new IP.'. The 'Replica IPs in DNS' section is expanded, showing a table of IP addresses. The first entry is '141.202.226.42', and the second is 'Click here to add new IP.'. The 'Switch Computer Name' option is set to 'Off'. The 'User-Defined Scripts' section is expanded, showing 'Is Alive', 'DB Management', and 'Action upon Success' options.

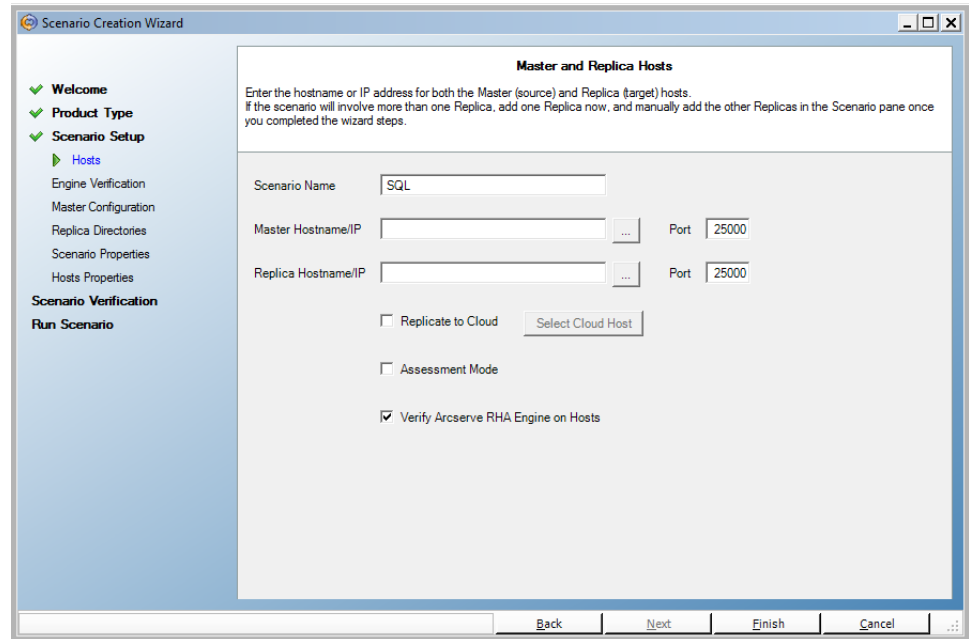
Property	Value
Move IP	Off
Redirect DNS	On
DNS Servers IPs	Off
DNS IP	141.202.226.10
DNS IP	Click here to add new IP.
DNS TTL (sec)	60
Active Directory Integrated	On
Master IPs in DNS	
IP Address	141.202.226.74
IP Address	Click here to add new IP.
Replica IPs in DNS	
IP Address	141.202.226.42
IP Address	Click here to add new IP.
Switch Computer Name	Off
User-Defined Scripts	
Is Alive	
DB Management	
Action upon Success	

After switchover, the Master's IP switches to the Replica host. This redirection method is applicable only when both Master and Replica host are on the same IP subnet.

Add Arcserve-IP to New Scenarios-SQL

Perform this procedure only if you are using the Move IP redirection method.

During the initial run of the Scenario Creation Wizard, enter the Arcserve-IP and Replica IP addresses instead of the server names.



Cluster Move IP

This section describes how you can configure the Move IP redirection method when working with clusters.

Note: If both Master *and* Replica are clusters, there are special configuration issues involved in the Move IP redirection process that are not detailed in this Guide. For a cluster-cluster scenario, use Redirect DNS or contact technical support to receive detailed instructions and guidance.

Use the Master Cluster

To use Move IP redirection with a clustered Master (MSCS with shared storage), you must add an additional IP resource to the Master Exchange resource group.

To use Cluster Move IP through the Master cluster

1. Open the Cluster Administrator.
2. In the Master Cluster Exchange Resource Group, create a new IP resource and name it **Arcserve-IP**.
3. Bring this resource online and verify that it is visible from the Replica via the ping command.

This new IP address is used for Arcserve RHA internal communication and replication. This is necessary since the current production IP address is not available on the Master cluster after switchover—it switches to the Replica server.

Use the Manager

This section details Cluster Move IP redirection using the Manager.

For New Scenarios

During the initial run of the Wizard, enter the Arcserve-IP and Replica IP addresses instead of the cluster virtual server names. The following screen shows the Arcserve-IP entered in the Master Hostname/IP field and the Replica Server IP address entered in the Replica Hostname/IP field.

The screenshot shows the 'Scenario Creation Wizard' window, specifically the 'Master and Replica Hosts' step. The window title is 'Scenario Creation Wizard'. On the left, there is a navigation pane with the following items: 'Welcome', 'Product Type', 'Scenario Setup' (expanded), 'Hosts' (selected), 'Engine Verification', 'Master Configuration', 'Replica Directories', 'Scenario Properties', 'Hosts Properties', 'Scenario Verification', and 'Run Scenario'. The main area is titled 'Master and Replica Hosts' and contains the following text: 'Enter the hostname or IP address for both the Master (source) and Replica (target) hosts. If the scenario will involve more than one Replica, add one Replica now, and manually add the other Replicas in the Scenario pane once you completed the wizard steps.' Below this text are the following fields and options:

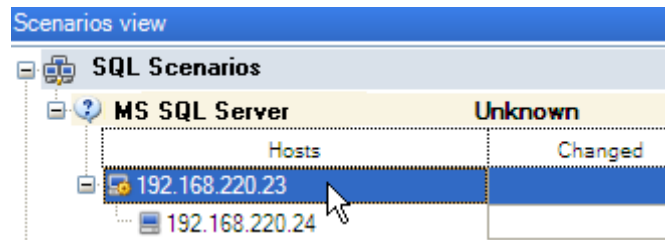
- Scenario Name:
- Master Hostname/IP: ... Port:
- Replica Hostname/IP: ... Port:
- Replicate to Cloud
- Assessment Mode
- Verify Arcserve RHA Engine on Hosts

At the bottom of the window, there are navigation buttons: 'Back', 'Next', 'Finish', 'Cancel', and a help icon.

For Existing Scenarios

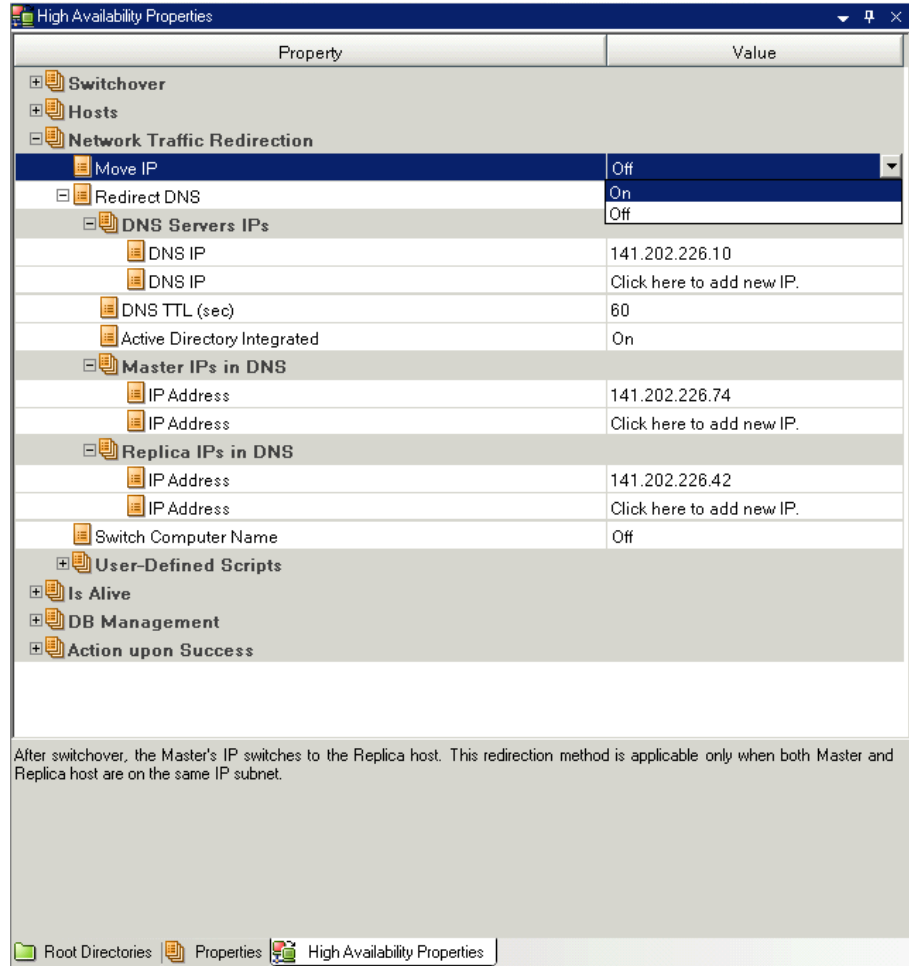
To use Cluster Move IP with existing scenarios

1. On the Scenario pane, select the required Master host.



2. Right-click the Master and select **Rename** from the pop-up menu. Then, enter the Arcserve-IP address.
3. On the Framework pane, select the **Switchover** tab and then select the Replica server as the switchover host.

- Set the **Move IP** option to On. Ensure that the IP address under **Move IP, IP/Mask** matches the production server IP address: this is the IP address that will switch over. If you are moving more than one IP address you can add multiple production IP addresses by selecting **Click here to add new IP/Mask**.



Switch Computer Name Redirection

If you are redirecting File Shares, in which clients connect via the Master server name, enable Switch Computer Name. For example, if the Master server name is fs01 and clients connect to \\fs01\sharename or \\fs01.domain.com\sharename, using the Switch Computer Name method redirects clients to the failover server. To use Switch Computer Name Redirection in the Active Directory environment, both Master and Replica must belong to the same domain.

It is also recommended to enable one other method. The most common method is to use both DNS Redirection and Switch Computer Name. Arcserve RHA makes the required computer name switch by assigning a temporary name to the Master server and taking over its computer name for use with the Replica server.

Arcserve RHA updates records directly and does not generally require a reboot. If, however, you encounter any problems after switchover, consider setting the reboot option to On and testing again.

Note: For Windows Server 2008 systems, you must reboot the computer after a switchover occurs when the switch computer name method is used. To do so, enable the property, Reboot after switchover. However, for Windows 2008 Cluster systems, reboot will not occur even if this property is enabled. You must reboot manually and ensure the SQL Server service is running.

Automatic Redirection Using Switch Computer Name

When possible during switchover, Arcserve RHA renames the master host to *masterhostname-RHA* and assigns its original name to the replica server. This step prevents name conflict, since the master's name is now assigned to the replica server. In this graceful case, if automatic reverse replication is set to on, Arcserve RHA starts the backward scenario automatically. If Automatic Reverse Replication is set to Off, run the scenario again manually by selecting the Run button or choosing Run from the Tools menu. Once the backward scenario has run and synchronization is complete, you can click Perform Switchover to switch back.

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.

Chapter 5: Managing Replication and High Availability Scenarios

This section contains the following topics:

[Scenario Properties](#) (see page 29)

[Run the Scenario from Outside the Wizard](#) (see page 32)

[View a Report](#) (see page 33)

[Stop a Scenario](#) (see page 34)

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. Select the scenario for which you wish to change properties, and then select the appropriate tab.

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.

When you select auto-discovery of database files, a SQL auto-discovery dialog opens. The dialog displays all the SQL instances and databases available on the master SQL host. This can be done in the wizard and/or in the scenario view. You can select the instances or databases that are needed.

Note: If you select an instance, you cannot deselect the master, model or msdb databases. You can only deselect user created databases.

There is also a checkbox called "Replicate new databases in the selected root directories". If checked, new databases that are created by the user after the scenario runs will be replicated.

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, Registry Synchronization or System State Protection) 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, choose email notification, or write to the event log.
- 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, and Port number of the Master and Replica
- Replication properties -- These properties differ for Master and Replica. See the Arcserve RHA Administration Guide for more information.
- Spool properties -- Set the size, minimum disk free size and directory path. See Spool Directory Settings for more information.
- Event notification properties -- specify a script to run, choose email notification, or write to event log.
- Report properties -- choose synchronization or replication reports, specify distribution or script execution
- (Replica) Scheduled Tasks -- set or suspend tasks, including Replica Integrity Testing for Assured Recovery. For more details, see the Arcserve RHA Administration Guide.
- (Replica) Recovery properties -- set delay, data rewind properties, or scheduled tasks for replica

Settings on the HA Properties Tab

These settings control how switchover and switchback are performed


- Switchover properties -- choose automatic or manual switchover, provide switchover hostname, and reverse replication settings
- Hosts properties -- specify the Master and Replica
- Network Traffic Redirection properties -- choose Move IP, Redirect DNS, Switch Computer Name 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

Run the Scenario from Outside the Wizard

After you create a scenario, you need to run it to start the replication process. Normally, before data changes on the Master will begin to be replicated on the Replica, the Master and the Replica need to be synchronized. Therefore, the first step in initiating a replication is synchronizing the Master and Replica servers. After the servers have been synchronized, online replication starts automatically, continuously updating the Replica with all of the changes that occur on the Master.

Note: In order for the replication process to succeed, verify that the user under which the Arcserve RHA Engine is running has Read permission on the Master, and Read and Write permissions on each replication root directory and included files, and on all participating Replica hosts.

To run the scenario from outside the wizard

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

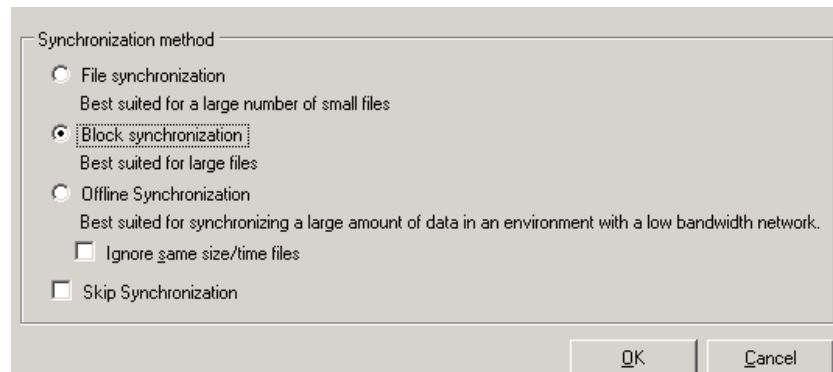
Before initiating synchronization and replication, Arcserve RHA verifies your scenario configuration. When verification completes successfully, Arcserve RHA Manager displays the message: *Are you sure you want to run scenario "scenario_name?"* If problems are discovered, the top pane displays any warning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

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\ARCserve RHA\Manager\reports`

To view a report

1. To view a report, first you need to open the Report Center. There are two ways to open it:
 - On the Overview Page, click the **Report Center** link on the **Quick Start** pane on the left.
2. From the **Tools** menu, select the **Reports** option and then **Show Scenario Reports**.

The Report Center opens in a new window.

The Report Center consists of two tables:

- The upper table - **Available Reports per Scenario** - contains a list of all scenarios that have reports, along 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.
3. 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)
XORMEXCHK7-1	Unknown	Today	03:29:37	Assured Recovery			811
XORMEXCHK7-1	Changes found	12/07/08	22:29:48	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.

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.

Chapter 6: 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 35)

[Initiate Switchover](#) (see page 37)

[Initiate Switchback](#) (see page 38)

[Switchover Considerations](#) (see page 41)

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. The following types of monitoring checks are available:

- Ping--a request sent to the Master to verify that the Master is up and responding
- Database check--a request that verifies the appropriate services are running and all databases are mounted
- 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 RHA 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.
- 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, except it 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. The timeout parameters are configurable and are more extensively covered in the *Arcserve RHA 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 has finished successfully.
- If you selected the Initiate Reverse Replication manually option, you must resynchronize data from Replica to Master, even after testing a clean switchover without a Master failure.


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

If you chose to initiate switchover automatically, after the Master is considered to be down, Arcserve HA automatically tries to restore the services and databases on it to their active state. First, Arcserve HA tries to restart the SQL services that it previously checked. If the services are running, it then tries to bring the relevant database's online. If all attempts failed, Arcserve HA initiates a switchover. These attempts to restore the services and databases are not performed if the switchover is initiated manually.

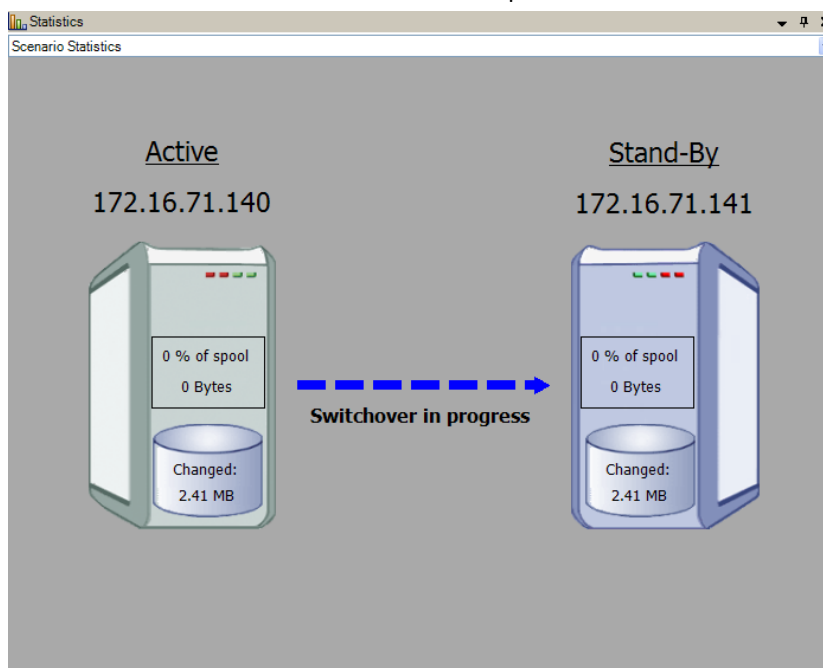
To initiate manual switchover

1. Open the Manager and select the desired scenario from the Scenario pane. Verify it is running.
2. Click on the **Perform Switchover** button, or select from the **Tools** menu the **Perform Switchover** option:



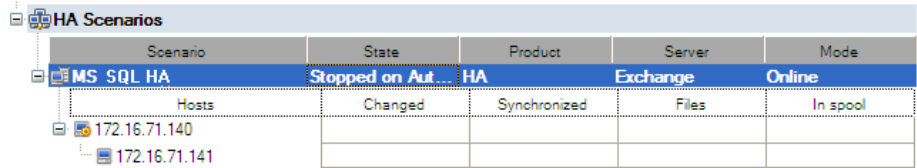
A confirmation message appears.

3. Click **OK** on the **Perform Switchover** confirmation message. This procedure initiates a switchover from the Master server to the Replica server:



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

- 4. After the switchover is completed the scenario stops:



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

Note: The only case in which the scenario may continue to run after switchover is when **automatic reverse replication** is defined as **Start automatically**.

In the Event pane a message appears, informing you that **Switchover completed**, and then that 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.


Note: The following steps are the same regardless of server type.

To initiate manual 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.

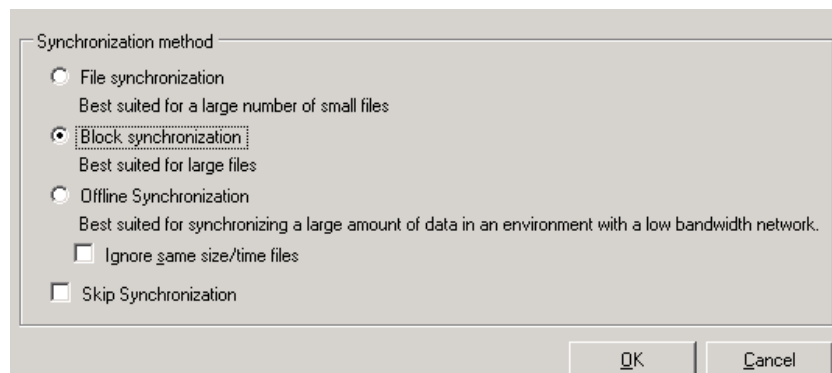
3. Perform one of the following:
 - If the scenario is already running, skip directly to Step 4
 - If the scenario is not running, perform these steps and then go to Step 4:
 - a. Click Run on the toolbar to start the scenario.

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.

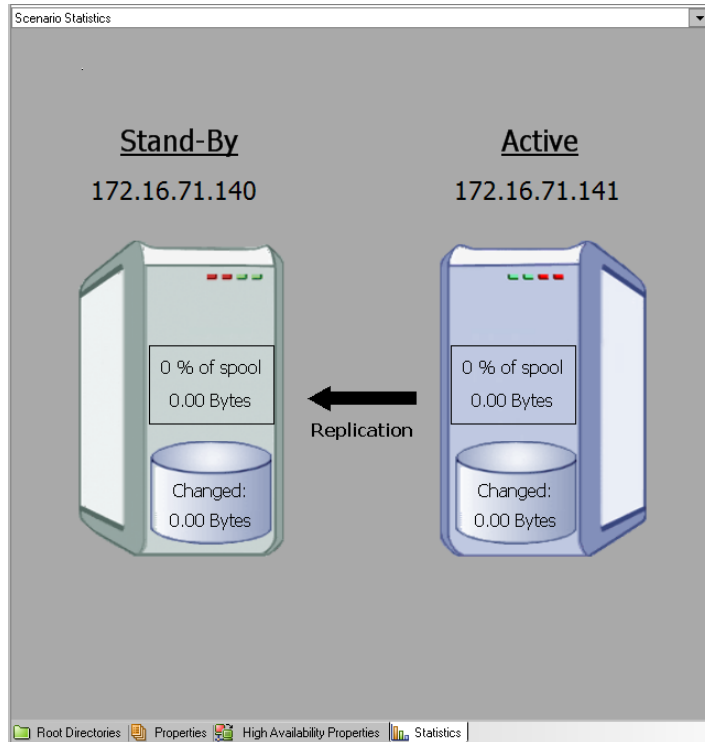


- b. Select a synchronization method from the Run dialog and click OK to start resynchronization.

Note: For SQL, 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 7: Recovering Data

This section contains the following topics:

[The Data Recovery Process](#) (see page 43)

[Recover Lost Data from Replica](#) (see page 43)

[Setting Bookmarks](#) (see page 45)

[Data Rewind](#) (see page 46)

[Recover Active Server](#) (see page 48)

The Data Recovery Process

When an event causes loss of Master data, the data can be restored from any Replica. The recovery process is in fact 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.
- Recovery 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 in order to initiate recovery.

Recover Lost Data from Replica

To recover all lost data from a Replica

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.



The **Restore Data** option is enabled.

4. From the **Tools** menu, select **Restore Data**, or click the **Restore Data** button on the Standard toolbar.

The **Recovery Method** page of the Restore Data wizard appears.

Note: If the **Data Rewind** property is set to On, another **Restore Data** dialog will appear. In this case, select the first option - Replace all data on Master with the data on Replica.

5. Click Next. The Synchronization Method page appears.
6. Make sure that the **Block Synchronization** method is selected, and click **Finish**.

Once you finished initiating the recovery process, Arcserve RHA builds a temporary reverse tree using the selected Replica as the root, and the Master as the terminating node. After the Master recovery process ends, the temporary scenario is deleted, and you receive the following message in the Event pane:

Synchronization finished.

7. By default, once a data recovery occurs a Synchronization Report is generated.

Now, the replication process can restart following the original scenario.

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.

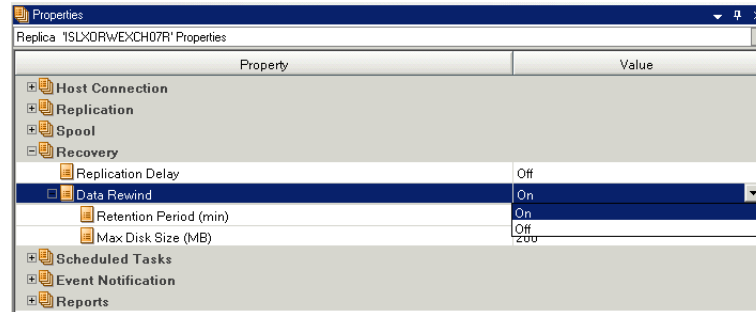
The bookmark is set.

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.



4. From the **Tools** menu, select **Restore Data**, or click the **Restore Data**  button.
The **Recovery Method** page of the Restore Data Wizard opens.

5. 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).

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.

6. Click **Next**. The **Rewind Point Selection** page is displayed.
7. 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 appropriate to the application you are protecting. These include modifications of folders and SQL checkpoints 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.

8. 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**.

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 RHA allows for this option through a process called Recover Active Server.

To use the Recover Active Server option, stop the scenario, and then 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 RHA 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 RHA 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, see [Recovering Servers](#) (see page 50).

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 occurs in a disaster situation 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 Settings](#) (see page 49)

[Rename the Microsoft SQL Server 2005](#) (see page 50)

[Recovering Servers](#) (see page 50)

Spool Settings

The Arcserve RHA spool is a folder on disk where data to be replicated is backed up (i.e., 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 Arcserve RHA 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 choose to change this value, it should be at least 10% of the total dataset size. For example, if you are replicating 50GB of data on a server you should ensure that at least 5GB of space is available for spool. Please note that this space is not pre-allocated.

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

Rename the Microsoft SQL Server 2005

If you run SQL Server 2005, the new name is recognized during the SQL service startup. There is no need to run Setup again or to reset the server name.

If there are any remote logins to the SQL Server, `sp_dropserver` may generate an error. To resolve the error, you may need to drop remote logins and rerun the procedure.

Note: In order to see the new SQL server in the Enterprise Manager, you need to delete an old server registration, and register the new name.

Recovering Servers

Arcserve RHA can detect when a Replica server is active and run the recovery process automatically. If recovery does not complete correctly for some reason, do the following:

- Perform the Recover Active Server procedure. For more information, see [Recover Active Server](#) (see page 48).
- If the Recover Active Server procedure does not resolve the issue, try one or more of the following manual tasks appropriate to the redirection method you use:
 - If IP Redirection is used, manually remove the IP. You cannot use this method for scenarios that do not support Move IP redirection (Hyper-V HA, CS HA). For more information, see [Manually Recover a Failed Server-Move IP Address](#). (see page 51)
 - If Switch Computer Name Redirection is used, manually switch the names. You cannot use this method for scenarios that do not support Switch Computer Name Redirection (Hyper-V HA, Exchange HA, vCenter HA if local Oracle is used). For more information, see [Manually Recover a Failed Server-Switch Computer Name](#) (see page 51).
 - If both IP and Switch Computer Name Redirection methods are used, manually remove the IP and switch the computer names. You cannot use this method for scenarios that do not support Move IP and Switch Computer Name redirection (Exchange, CS HA). For more information, refer to the topic, [Manually Recover a Failed Server-IP and Switch Computer Name](#) (see page 53).

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.

Manually Recover a Failed Server-Switch Computer Name

If Switch Computer Name Redirection is used, you must switch the names manually. You cannot use this method for scenarios that do not support Switch Computer Name Redirection (Hyper-V HA, Exchange HA, vCenter HA if local Oracle is used).

To manually recover a failed server using the Switch Computer Name redirection method

1. Boot the Master server without a network connection, to avoid duplicate network names.
2. Rename the server to <NewServerName>-RHA and move it to a temporary workgroup.

For example, if the server is "Server1", rename it to "Server1-RHA". You will be required to reboot this computer. After the reboot completes, the following error appears: "At least one Service could not be started." Ignore this, it is normal under these circumstances because the Arcserve RHA Engine usually runs in a domain account.

3. Connect to the network.
4. Rejoin the domain, make sure that you use the -RHA name assigned in step 2.
5. Reboot the computer.
6. 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 now standby.)
7. Wait for synchronization to complete. Perform a manual switchover to make the Master server active. It is recommended that you do so outside of normal business hours.

Manually Recover Failed Server-IP and Switch Computer Name

If both IP and Switch Computer Name Redirection methods are used, you must remove the IP address and switch the computer names manually. You cannot use this method for scenarios that do not support Move IP Address and Switch Computer Name redirection (Exchange, CS HA).

To manually recover a failed server using both IP and Switch Computer Name redirection methods

1. Repair any hardware problems that could have caused the switchover, if any.
2. Reboot the server without a network connection to prevent IP conflicts.
3. From the TCP/IP properties dialog, remove the additional IP address.
4. From the System Properties, Computer Name dialog, change the Computer Name to <ServerName>-RHA. For example, if your server is called Server 3, rename it to Server 3-RHA.
5. Assign the server to a temporary workgroup.
6. Restart the computer to enable your changes to take effect. When rebooting completes, reconnect to the network now. Ignore the message, "At least one service failed during system startup." This is normal because the Engine runs in a domain, which is not currently available.
7. Rejoin the domain, making sure you use the -RHA name, and reboot again.
8. The reverse scenario begins and the Replica server assumes the active role. Wait while synchronization completes.
9. Perform a manual switchover by clicking the Perform Switchover button from the toolbar, to return the active role to the Master server.

Index

A

- About Clusters • 11
- About This Guide • 7
- Add Arcserve-IP to Existing Scenarios-SQL • 22
- Add Arcserve-IP to New Scenarios-SQL • 24
- Add IP on the Master Server • 20
- Additional Information and Tips • 49
- Arcserve Product References • 3
- Automatic Redirection Using Switch Computer Name • 28

B

- Base Configuration • 9

C

- Cluster Move IP • 11, 24
- Contact Arcserve • 3
- Create SQL Server HA Scenarios • 14
- Create SQL Server Replication Scenarios • 13
- Creating and Using a High Availability Scenario • 13
- Creating Replication and High Availability Scenarios • 13

D

- Data Rewind • 46
- DNS Redirection • 19
- Documentation Changes • 4

F

- For Existing Scenarios • 26
- For New Scenarios • 25

H

- How Redirection Works • 19
- How Switchover and Switchback Work • 35

I

- Initiate Switchback • 38
- Initiate Switchover • 37
- Introduction • 7

L

- Log On Account Conditions • 10

M

- Manage Services • 16
- Managing Replication and High Availability Scenarios • 29
- Manually Recover a Failed Server-Move IP Address • 50, 51
- Manually Recover a Failed Server-Switch Computer Name • 52
- Manually Recover Failed Server-IP and Switch Computer Name • 50, 53
- Move IP Redirection • 20

P

- Protecting Microsoft SQL Server Environments • 9

R

- Recover Active Server • 48, 50
- Recover Lost Data from Replica • 43
- Recovering Data • 43
- Recovering Servers • 48, 50
- Redirection Methods • 14, 15, 19
- Register Arcserve RHA Licenses • 11
- Related Documentation • 7
- Rename the Microsoft SQL Server 2005 • 50
- Run the Scenario from Outside the Wizard • 32

S

- Scenario Properties • 29
- Scripts Redirection • 28
- Server Requirements • 9
- Setting Bookmarks • 45
- Spool Settings • 49
- SQL Server Configuration Requirements • 10
- SQL Servers Operating in a Workgroup • 11
- Stop a Scenario • 14, 15, 34
- Switch Computer Name Redirection • 27
- Switching Over and Switching Back • 15, 35
- Switchover Considerations • 41

T

The Data Recovery Process • 43

U

Use the Manager • 25

Use the Master Cluster • 25

V

View a Report • 33