# Administration Guide

**Arcserve® Replication and High Availability**

**r16.5**

arcserve®

# Legal Notices

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.

© 2017 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.

Providing Feedback About Product Documentation

If you have comments or questions about Arcserve product documentation, please contact us.

# Contents

# Chapter 5: Running the Replication Process ................................. 143

# Chapter 1: Introduction

This section contains general information about the Arcserve Replication and High Availability (Arcserve RHA) products, and their various modules. It briefly lists new features, describes how replication and high availability work, and how the various modules function in the replication process.

This section contains the following topics:

# About this Guide

This guide contains all of the necessary information for configuring and running Arcserve RHA. It describes and provides instructions on how to perform the following procedures:

- synchronizing

- replicating and recovering data

- monitoring procedures

- generating reports

- switching over from the production server to the Replica standby server, and switching back

- protecting the Control Service

**Important!** This guide applies to replication, high availability and assured recovery products. Throughout this document, the term, Arcserve RHA refers to all products, unless otherwise specified.

This guide focuses on the generic **File Server** replication and high availability solutions, but it also provides information about other application and database servers and high availability solutions.

For more detailed instructions involving scenarios tailored to specific applications such as Microsoft Exchange or SQL Servers, see the appropriate Operation Guide. You can find the most up-to-date Operation Guides for each application on the Arcserve Support site. For more information about viewing application-specific Operation Guides, see Related Documentation.

# View Related Documentation

The *Arcserve RHA Administration Guide* is meant to be used with the following sup-plemental guides.

- *Arcserve RHA Installation Guide - contains information related to installing and configuring Arcserve RHA*

- *Arcserve RHA PowerShell Commands Guide - contains command line reference information*

In addition, Operation Guides provide the details, examples and settings necessary for successfully using the software in specific application or database server envir-onments. Replication and High Availability information is provided in these guides. (In previous versions of Arcserve RHA, separate Replication (Disaster Recovery) and High Availability (HA) guides were provided.)

- *Arcserve RHA for Microsoft SQL Server Operation Guide*

- *Arcserve RHA for Microsoft Exchange Server Operation Guide*

- *Arcserve RHA for Microsoft SharePoint Server Operation Guide*

- *Arcserve RHA for Microsoft Dynamics CRM Operation Guide*

- *Arcserve RHA for Oracle Server Operation Guide*

- *Arcserve RHA for Microsoft IIS Server Operation Guide*

- *Arcserve RHA for Blackberry Enterprise Server Operation Guide*

- *Arcserve RHA for UNIX and Linux Operation Guide*

- *Arcserve RHA for Virtualized Server Environments Operation Guide*

For information on integrating Arcserve RHA with Arcserve Backup, see the *Arc-serve Backup Arcserve RHA Integration Guide,* in the Arcserve Backup doc-umentation set.

Arcserve Replication and High Availability documentation contains specific guides and release notes for all major releases and service packs. Click links below to access documentation.

- [Arcserve RHA r16.5 Release Notes](#)
- [Arcserve RHA r16.5 Bookshelf](#)

# Supported Application and Database Servers

Replication and high availability capabilities are custom-tailored for the following application and database servers, for both 32-bit and 64-bit Windows:

- Microsoft File Server -- protection for virtually any application or data type, including databases, so long as the OS platform and file system are supported.

- Microsoft Exchange Server -- protection for your Exchange servers

- Microsoft SharePoint Server -- protection for your SharePoint servers

- Microsoft SQL Server -- protection for SQL servers, including the BlackBerry Enterprise Server database.

- Microsoft IIS Server -- protection for Internet Information Systems.

- Microsoft Hyper-V -- protection for virtualized environments.

- Microsoft Dynamics CRM -- protection for your Dynamics CRM servers.

- Oracle databases -- protection for Oracle databases.

- VMware vCenter Server -- protection for virtualized environments.

- Arcserve RHA Control Service -- protection for this Arcserve RHA component.

- Full System HA - transfer of an entire computer (physical or virtual) to a virtual machine (Hyper-V, VMware ESX, and Citrix XEN Server).

**Note:** BlackBerry Enterprise Servers can be protected using Arcserve RHA for Microsoft SQL Server or File Server, depending upon your configuration. For details, see the Arcserve RHA for Windows BlackBerry Enterprise Server Operation Guide.

For an up-to-date list of supported platforms and applications, see the *Supported Configurations* article on Arcserve.com.

The properties you configure during scenario creation vary depending on the application or database server you are protecting. This *Administration Guide* provides descriptions of each property; for specific scenario creation instructions, refer to the appropriate *Operation Guide.*

# Arcserve RHA Concepts

The following concepts describe how Arcserve RHA protects your server environments.

[How Synchronization Works](#)

[How Replication Works](#)

[How Recovery Works](#)

[How Data Rewind Works](#)

[How Replication Suspension Works](#)

[How High Availability Works](#)

# How Synchronization Works

Synchronization of files is the process of making the set of files to be protected identical on the Master and Replica servers. It is usually necessary to synchronize the Master and Replica as the initial step of a replication scenario.

**Important!** We strongly recommend running the initial synchronization during off-peak hours.

Files identified as sparse are now synchronized as such if the volume on the Replica server supports them. If not, sparse file attributes are lost during the replication or synchronization processes.

This section contains the following topics:

- Synchronization Methods
- Synchronization Filters
- Automatic Synchronization
- Simultaneous Synchronization and Replication
- Reporting Synchronization Differences

# Synchronization Methods

To properly synchronize the Master and the Replica, it is first necessary to perform a comparison of their two file structures. This comparison determines what content (files and folders) on the Master is missing or different from the content on the Replica.You have two ways to start Synchronization:

- Click the Synchronize button on the Manager toolbar
- Click the Run button on the Manager toolbar

Select a synchronization mode, each with a different comparison algorithm and method of operation:

**File synchronization**

In file synchronization, the Replica server sends its snapshot to the Master server, which uses it to drive the information and content for the data comparison. After the comparison is performed, the Engine on the Master server sends a sequence of commands to the Replica. These commands:

- Delete files that exist only on the target
- List the entire file contents of files that either exist only on the Master, or that exist on the Replica but differ from the version on the Master

This method is best suited for File Servers or application servers with many relatively small files.

**Note**: When performing File Synchronization, Arcserve RHA does not update the data transferred percentage until the file transfer is completed.

**Block synchronization**

In block synchronization, the Engine performs a block-by-block comparison of the Master and Replica files, and copies over only those blocks that are different. When differences exist between files, instead of requiring the transfer of the entire file, block synchronization transfers changes only.

This method is best suited for database applications, such as Microsoft Exchange Server, Oracle, and SQL Server, or application servers with large files.

**Offline synchronization (Available only from the Run dialog)**

In offline synchronization, data is copied from the Master server to an external device, and from the external device to the Replica server. Full System scenario

supports Offline synchronization method for forward, backward, and BMR restore scenario.

This method is best suited for transferring huge data volumes without much impact from a low network bandwidth. This option is available only when running a scenario and does not apply to scenarios with scheduled replication, or scenarios in which the Master is a UNIX/Linux host. For more information, see How Offline Synchronization Works.

**Volume Synchronization (Available only for Full System Scenarios)**

In volume synchronization, the Master server copies the volume as whole instead of a file or a block. All volumes on a disk are synchronized at a volume level. RHA accesses the bitmap and only copies the used data. For example, when the used volume is 25 GB of a 40 GB volume, only the used 25 GB is copied.

**Note:** If the size of volumes on the Master is larger than the size on Replica, RHA switches to block synchronization.

The comparison can be configured to consider only file size and modification time to determine whether two files differ, or it can perform a check of the actual contents of the data. The former approach, which is not valid in the case of database applications, can be a legitimate way to increase the comparison process on a File Server scenario.

You can Filter or Skip synchronization.

# Synchronization Filters

Before the synchronization starts, you can filter the synchronization process. The filter, called **Ignore files of same size/time,** appears on the **Run** dialog, and it can be either enabled or disabled.

The **Ignore files of same size/time** filter allows the data comparison between the Master and the Replica to consider only file size and modification time, when determining whether two files differ. It skips comparison of files with the same path, name, size and modification time, under the assumption that these files are identical. This approach is not valid in the case of database applications, but it can be an efficient way to significantly speed up the comparison process for a File server solution, and reduce the overall synchronization time dramatically.

**Important!** Do *not* skip synchronization unless you are absolutely sure that files on the Master and Replica are identical.

# Automatic Synchronization

You can configure your system to perform automatic synchronization after certain events occur. The Automatic Synchronization property causes the Master and Replica servers to be resynchronized automatically if one of the following events occurs:

- The Replica is restarted
- The Master is restarted

**Note**: If the Master spool overflows because of network failure, the servers are resynchronized automatically when the connection is restored

You can set the Automatic Synchronization property in the Scenario Properties, Replication Group.

# Simultaneous Synchronization and Replication

Simultaneous Synchronization and Replication means that servers can be synchronized while files are in active use and being updated. All changes that occur while the initial synchronization is performed are replicated without any need for administrative intervention.

# Reporting Synchronization Differences

The data sets on the Master and Replica servers can be checked for differences without performing resynchronization using the Difference Report option.

# How Replication Works

The replication mechanism maintains identical copies of files and databases on the Master and Replica. This replication is done by real-time capture of byte-level changes in files on the Master server, using a file-system filter-driver. The captured changes are transmitted asynchronously to the Replica servers using the Engine. The replication process does not interfere with write operations.

To accommodate real-time replication of all types of files, the following replication modes are supported:

- **Online mode** - Replicates captured changes of files, even if files are always open (as is the case for most database and mail servers). This mode maintains the order of file system operations. In this mode, the Engine records all I/O operations related to the root directories in journal files. The journal files are then sent to the Replicas where the operations that were recorded in the journal are replayed on the replicated files.

- **Scheduled mode** - Synchronization of servers takes place at fixed times. There is no online replication in this mode, however, online changes made during synchronization are replicated. (Scheduled replication cannot be performed using Offline Synchronization.)

You can assess the accurate bandwidth usage and compression ratio benchmarking that is needed for replication, without actually replicating data. When you select the Assessment mode, no replication occurs but statistics are gathered. A report is provided after the assessment process is completed.

Sparse files are now supported. Sparse files are typically very large files that contain mostly zeros. When NTFS file systems encounter large runs of zero data, they do not explicitly write the zeros to disk. Instead, the file system maintains a reference that tracks the locations of these zero runs. Though the file size is still reported as usual, much less disk space is consumed. Arcserve RHA ensures consistency of content in sparse files. You cannot replicate sparse files to a Replica server that does not support them, such as a FAT32 Replica.

Sparse file operations are transparent; they are handled internally.

# How Recovery Works

When Master server data is lost or corrupted for any reason, you can recover the data from any of the Replica servers participating in the scenario. The Restore Data option activates a synchronization process in the reverse direction: from Replica to Master.

When recovery is initiated, the Manager builds a temporary tree, containing a single branch. In this tree, the Replica becomes the source of data and the original Master becomes the target (the terminating Replica). After the synchronization process completes, the Manager reverts to the original replication scenario, and continues working.

**Important!** All file system activity must be halted on the Master host until the recovery process finishes.

# How Data Rewind Works

Data rewind is technology that lets you repair a corrupt file by rewinding it back in time as if it were a tape. Because replication continuously updates source data to another machine, a Replica always holds the same data as in the Master. In the case of data corruption, recovering files from the Replica will not help, since chances are high that data on the Replica is also corrupted.

Data Rewind can be compared to the Undo feature of various productivity applications where user actions can be cancelled, which brings the file to a previous state in time. Data Rewind is based on rewind journals that store I/O operation information that result in modified files. Using the rewind journal, it is possible to *undo* I/O operations, and rewind the file to a previous point in time, to a valid, non-corrupted state.

# How Replication Suspension Works

At times it may be necessary to suspend updates on a Replica machine in order to perform system maintenance or some other form of processing that does not modify the replicated data there. It is not desirable to stop replication since this requires a resynchronization afterward.

The replication suspension feature makes this possible. Replication may be suspended either manually or on a scheduled basis. During the suspension period, all changes are spooled on the Master or on the Replica located upstream of the suspended Replica. In other words, changes continue to be recorded for update on the suspended Replica, but are not actually transferred until replication is resumed. After replication resumes, the accumulated changes are transferred and applied without the need to perform a resynchronization of the data.

To suspend replication, choose Tools, Suspend Replication from the Manager menu.

# How High Availability Works

Arcserve RHA monitors all critical events, including global server failure and all database service failures. You may configure Arcserve RHA to automatically reverse server roles, called failover, or manually switch server roles, called switchover, when a critical event is detected. This means the Replica server becomes active while the Master server stands by.

**Note**: Automatic failover and manual switchover settings are configured during High Availability scenario creation. You may also configure these settings from the Arcserve RHA Manager Properties tab and expand the Switchover Settings properties group.

When the Master server becomes unavailable, its activities are failed over automatically to a remote site (Replica). The failover, which is transparent to the user, includes immediate startup of a synchronized standby database. All users are redirected to it in minimum time. All this is done without any need to reconfigure either clients or the network.

Redirection can be based on the following methods:

- Move IP (if the standby site is implemented within the same network segment)

- Redirect DNS, can be used on a local network or when the remote standby site is located on a different IP network (cross-network switchover)

- Switch the server hostname/NetBIOS name

**Note**: You can also apply user-defined scripts that add or replace the built-in redirection methods. Identify Network Traffic Direction scripts are required to fully support custom, or customized, redirection methods. Custom scripts or batch files are used to identify the active server. This script determines if the forward or backward scenario will run when the scenario is started. The script runs on both the Master and Replica: the one that returns zero is active. If both return zero, a conflict is reported.

Redirection method selection is based on the requirements of the application being protected; certain methods may not apply to a particular scenario. For more information, see the application-specific Operation Guide.

## How File Server Assured Recovery Works

When creating File Server scenarios, you may elect to choose the option, Integrity Testing for Assured Recovery (AR) from the Select Server and Product Type dialog. You can set a schedule for testing. Double-click the value for Schedule to open the Assured Recovery hours screen.

By default, File Server AR takes VSS snapshots on the Replica during the test. These snapshots use volume space on the Replica. VSS Snapshots are off by default to prevent disk space issues.

**Note:** Since there is no application in a File Server scenario, Assured Recovery testing requires custom scripts.

# Limitations

Only one-way, asynchronous, replication is permitted, and the Replica database must be offline. Bidirectional replication is not supported. However, cross replication with different data sets is supported. A server running Arcserve RHA can act as a Master and Replica for an unlimited number of scenarios so long as each data set only has a single Master server, i.e., one way replication.

# Replication and High Availability Components

Arcserve RHA consists of the following components:

- Control Service

- Engine

- Management Center - consists of three components: Overview Page, Manager, and Report Center.

- PowerShell

- Arcserve RHA for Microsoft Failover Cluster

# Control Service

The Control Service functions as the single-point-of-control in the Replication or High Availability operation. It contains the entire dataset of the existing scenarios. The Control Service communicates with the Engines and the Managers. It is responsible for the management of all scenario-related-tasks, such as, creation, configuration, monitoring, and running of the scenarios.

The Control Service receives requests from the Managers, processes them, converts them to particular commands, and passes them on to the Engines. Then, the Control Service receives up-to-date data and events from the Engines, and sends back information and statistics about the scenario's state to the Manager.

The Control Service is also responsible for the authentication and authorization of users. It can also serve as a central point for report handling and storage. The information and statistics that are accumulated by the Control Service can be presented to the user through the Overview Page, Manager, Report Center and PowerShell.

All the scenario files are kept on the server that runs the Control Service. If the Control Service is down, the scenario functioning will not be affected. However, for receiving information about the scenario's state, the Control Service must be active. For best results, install the Control Service on a standalone host. If this is not possible, you can install the Control Service on either the Master or Replica servers. However, if the server is down, the connection with the Control Service is lost and scenarios will be unmanageable.

You may protect the Arcserve RHA Control Service in separate scenarios. For more information, see Protecting the Control Service in the *Arcserve RHA Administration Guide*.

# Engine

The Engine is a service that must be running before any scenario can start. It is installed on every server participating in any given scenario, meaning the Master (source) and Replica (target) hosts. Each Engine supports both Master and Replica functionality, for both Replication and High Availability scenarios. It may participate in multiple scenarios and serve in a different role in each scenario. Engines can be installed either locally on each host at a time, or through a remote installer on numerous hosts at once, and can be installed during scenario creation, if needed.

# Management Center

The Management Center consists of three components, none of which requires any manual installation:

- **Overview Page** - a statistical overview of the Replication and High Availability scenario state.

- **Manager** - a User Interface that lets you create, configure, manage and monitor scenarios. This is a GUI application that is activated from the Overview Page by clicking the Scenario Management link.



- **Report Center** - an interface that gathers all existing reports, along with information about the available reports per scenario. You can decide where these reports will be stored, and for how long they will be displayed and saved in the Report Center.

# PowerShell

The PowerShell is offered as an alternative if you do not want to manage the replication process using the Manager graphic user interface. It enlarges and facilitates the capabilities of the CLI provided in previous versions, and it supports both replication and HA operations.

The PowerShell is a command-line shell and scripting environment that allows you to configure a replication scenario and control and monitor the replication process. All scenarios managed by the PowerShell look and operate exactly as the ones that are managed by the Manager, and they are automatically saved in the same default location: INSTALL_DIR/ws_scenarios.

The PowerShell is based on the standard Windows PowerShell™, which comes with a large set of built-in commands with a consistent interface. The PowerShell component adds a number of scenario-related-commands, called snap-ins, which facilitate scenario management.

# Arcserve RHA for Microsoft Failover Cluster

Arcserve RHA for Microsoft Failover Cluster includes a virtual disk resource plug-in and a lightweight interface installed on every node in the cluster.

Though high availability is built in to Microsoft cluster environments, shared storage could still be a single point of failure. Arcserve RHA protects cluster storage, allowing you to replicate data to a disk resource on or off-premises.

# How to Deploy Arcserve RHA

Deploying Arcserve RHA components depends on the size of your IT enterprise network and your replication and HA needs. However, there are certain guidelines that you should follow when designing your environment and deploying different components on a Windows platform. For information regarding an efficient deployment, see the *Arcserve RHAInstallation Guide*.

In general, you install the Engine on pairs of servers - a Master and a Replica. The Control Service should be installed on a standalone server so that it can be protected in its own scenario. For more information, see the topic, Protecting the Control Service.

# Chapter 2: Exploring the Manager

This section introduces the Manager, its components and functionality. It explains how to log in to the Management Center and Manager, and describes the structure, menus, buttons and functions available in the Manager main window.

This section contains the following topics:

# Log Into the Management Center

The Management Center and Manager do not require any component or application installed in advance. It is based on a one-click-installation procedure that can be performed from any workstation that has a network connection and a Web browser. To log in, you will need the following:

- The Hostname/IP Address and Port Number of the server where the Control Service is installed.

- Your User Name, Password and Domain for that host.

**To open the Manager**

1. Open Internet Explorer. On the **Address** box, enter the Control Service Host Name/IP Address and Port Number as follows:

   http://host_name:port_no/start_page.aspx

   The **Login** dialog appears.

   Be aware of the following:

   - If you are opening the Management Center from the machine where the Control Service is installed, you can use the default parameters:

     http://localhost:8088/start_page.aspx

   - If you selected the **SSL Configuration** option during the installation of the Control Service**,** when you open the Overview page, you need to use the hostname of the Control Service machine (instead of its IP Address). Enter the Control Service Host Name and Port No. as follows:

     https://host_name:port_no/start_page.aspx

2. Enter a User Name, Password and Domain and click **Log In.**

   **Important**! To log in to the Management Center, you must be a member of the Administrators Group on the Local machine where the Control Service is installed.

   The **Overview page** appears.

3. On the **Quick Start** toolbar on left, click the **Scenario Management** option.

   A progress bar appears, indicating that the Manager component is currently installed on the local machine.

4. After the Manager installation is completed, the Manager appears.

   **Important!** Multiple administrators can simultaneously access the Manager, and they can make any changes anytime they need, depending on their

privileges. The last update will be effective as the latest state of the scenario. Therefore, when multiple administrators are working with the Manager on the same time, it is important to be aware that one administrator can unintentionally overwrite the changes another administrator just did. We recommend taking internal measures to prevent the occurrence of this event.

# Support and Community Access Links

The Overview screen has been redesigned to incorporate the latest Internet technology. Across the top of the Overview screen, an RSS feed cycles through product headlines. Click a headline to view the whole story posted at the Expert Advice Center.

The Support and Community Access section of the Arcserve RHA home page provides links to various product support sites.
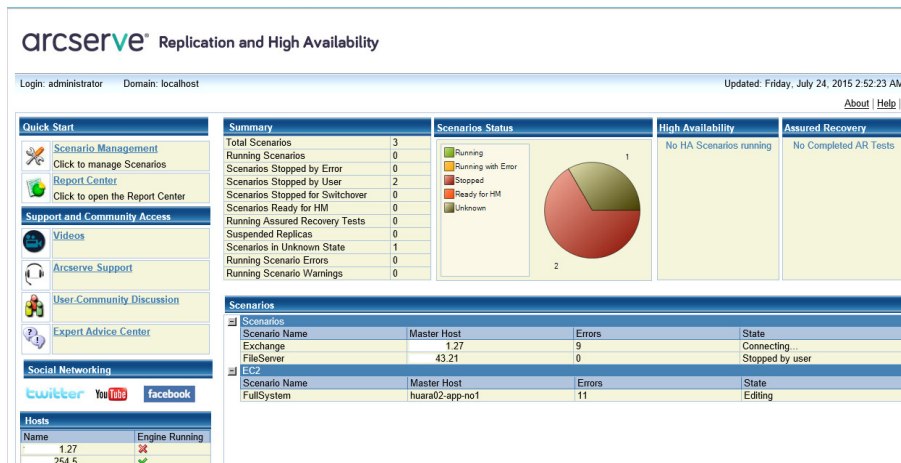


**Videos**

Click here to access available Arcserve RHA how-to videos for basic procedures. You can also view videos directly from YouTube. (Click the YouTube link in the Social Networking panel.)

**Arcserve Support**

Click here to go to the One Stop Support site where you can resolve issues and access important product information.

**Provide Your Feedback**

Submit questions and share your ideas for new product features via our *GetSatisfaction* partner. All information is submitted directly to the product development team.

**User-Community Discussion**

Be part of the Arcserve RHA User Community. Click here to share your tips and best practices or ask questions.

**Expert Advice Center**

Subscribe to the news feed to get the latest product news and information as well as link to other Arcserve RHA-related information.

**Social Network Panel**

You can even follow us on Twitter or friend us on Facebook to keep pace with product changes. Click the YouTube link to access videos.

Click All Feeds (located in the top right corner) to subscribe to any or all of the sites. When a site is updated, you are notified.

Feeds (disabled by default) and social networking links (enabled by default) and can be enabled or disabled as needed. Edit the configuration file as follows to enable these settings:

1. Open the web.config file located at [Install dir]/ws_root.

2. Locate and set the following settings:

   <appSettings>

   <add key="SocialNetwork_Visible" value="true" />

   <add key="GoogleFeed_Visible" value="false" />

   </appSettings>

3. Save the web.config file.

# Explore the Arcserve RHA Manager Screen

After logging in to the application, the Manager is displayed, enabling you to access all the Manager menus, toolbar functions and panes.

Unless a scenario exists, most of the user areas are blank. If active scenarios exist, they are shown on the left side of the Manager screen.
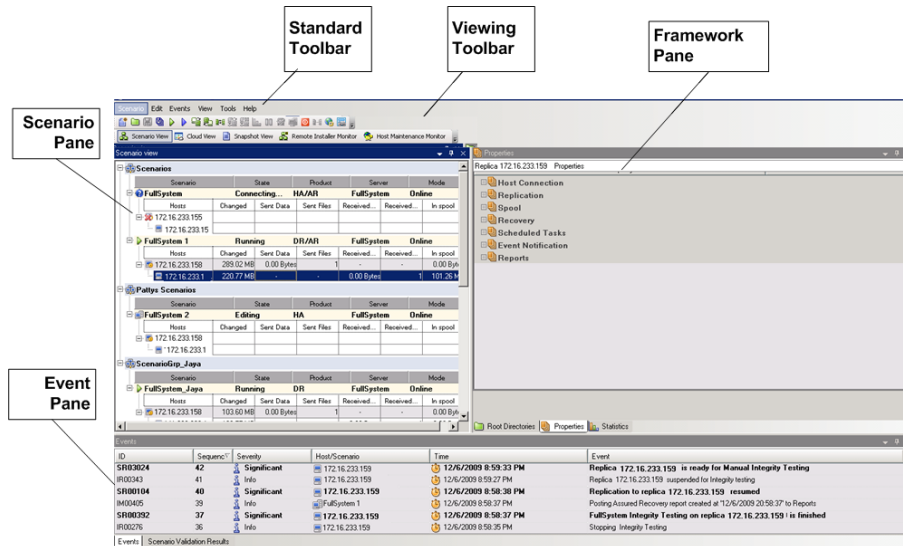


**Note**: Some of the panes and options are visible and enabled only with the appropriate product license.

# Manager Screen Panes

The Manager screen is divided into several areas:

- The application's name and the Control Service's connection details appear in the top left corner of the title bar; beneath it appear the menu line, the Standard toolbar and the Viewing toolbar.

- The Scenario pane appears on the left. In this pane, the existing scenarios, including their replication trees, are displayed.

- The Framework pane appears on the right. In this pane, configurable property lists are displayed - the Scenario, Master, Replica, HA and Template properties. The displayed list depends on the item selected in the Scenario pane, or on the item chosen from the pane's drop-down list. The Framework pane also shows two, three or four tabs, depending on the selected solution and on the scenario state. These tabs include:

  - Root Directories

  - Properties

  - High Availability Properties

  - Statistics

  The Properties displayed on each tab are described more fully in the topic, Setting Scenario Properties.

- The Events pane is below the screen's dividing line.

**Note**: The actual placement of the panes can vary, since they can be moved and resized. Also, you can hide panes, including the toolbar and status bar according to the selections in the View menu.

# Viewing and Arranging the Manager Screen

Arcserve RHA offers you different ways to get a good view of your work, depending on the task at hand. You can arrange your workspace according to your current needs.

**Note**: The last view setting you use is saved for the next session.

The following topics describe ways to view the Manager Screen:

- Viewing Options
- Customize the Scenario View
- Rearrange Panes

# Viewing Options

The **View** menu contains all existing views. Selecting a view option from the menu will either open or close the indicated pane.

**To reset your Manager layout**

1.  From the **View** menu, select the Reset option.

The original view settings are restored.

# Customize the Scenario View

The Scenario pane shows the current status for all scenarios in one pane, enabling you to monitor multiple scenarios at once. You can customize the way the scenario information columns are displayed.

**To customize your scenario view**

1. From the View menu, select the Customize Scenario View option.

   The Customize scenario view dialog appears.

2. Select the fields you want to display in the Scenario pane and click OK.

   **Note**: The Started by field indicate the user that initiated the Run for the specific scenario.

   The fields you selected appear as columns in the Scenario pane.

# Rearrange Panes

You can dock, stack, hide, show and float the Manager panes according to your needs.

Dock a Pane

Stack Panes

Hide a Pane

# Dock a Pane

The Pane Docking tool, or guide diamond, is a built-in feature that appears automatically whenever you move a pane.

**To dock a pane**

1. From the Manager screen, click a pane's title bar and start dragging it. The Pane Docking tool appears.
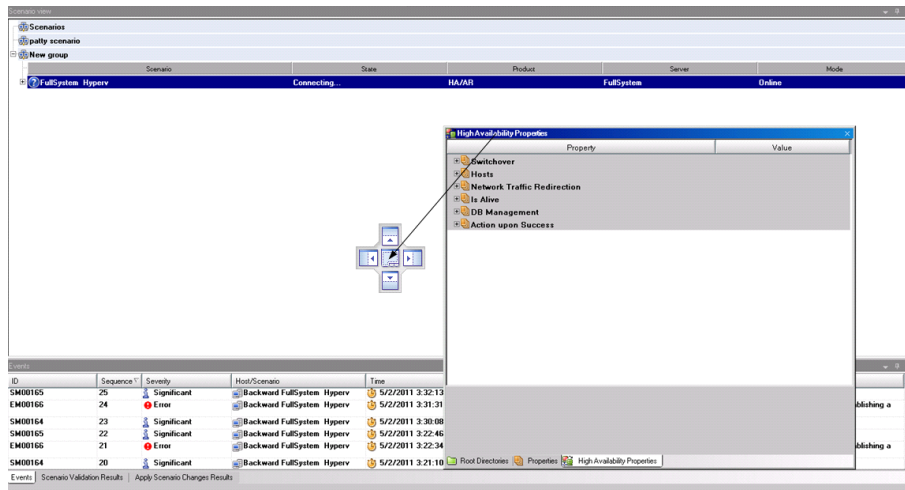


2. Move your pointer over the corresponding portion of the Pane Docking tool. You can also use the arrows at the edges of the window.

3. Release the mouse button when the pane reaches the desired position.

The pane is docked in the new screen location.

# Stack Panes

Stacking panes on top of each other presents each as a tab in the Manager screen.

**To stack panes**

1.  From the Manager screen, click a pane's title bar and start dragging it. The Pane Docking tool appears.
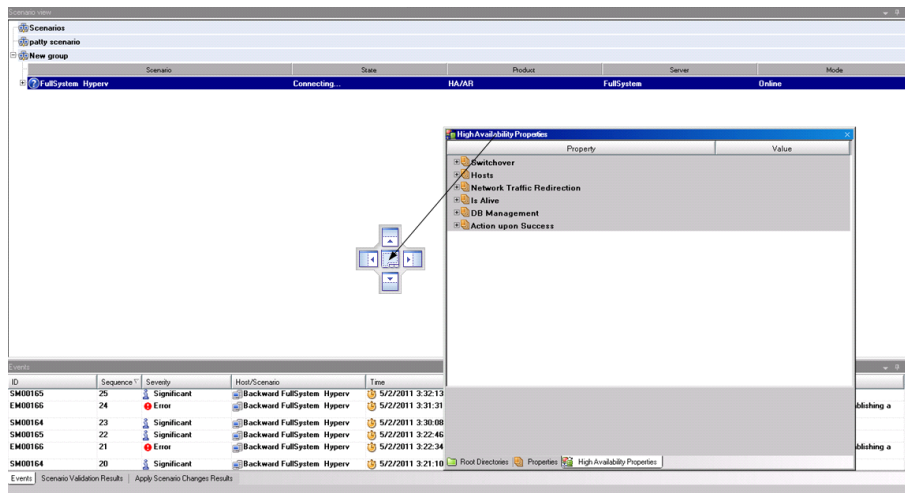


2.  When the pane you are dragging reaches a docked pane on which you want to stack it, move the pointer to the center of the guide diamond.

3.  Release the mouse button.

The pane is now accessible by clicking its tab.

# Hide a Pane

You can completely hide a pane or hide it only while working on another pane. You can return to the hidden pane by clicking its tab.

Restore hidden panes by clicking the View, Reset menu option.

**To hide a pane**

1. From the Manager screen, right-click a pane's title bar. A short-cut menu appears.
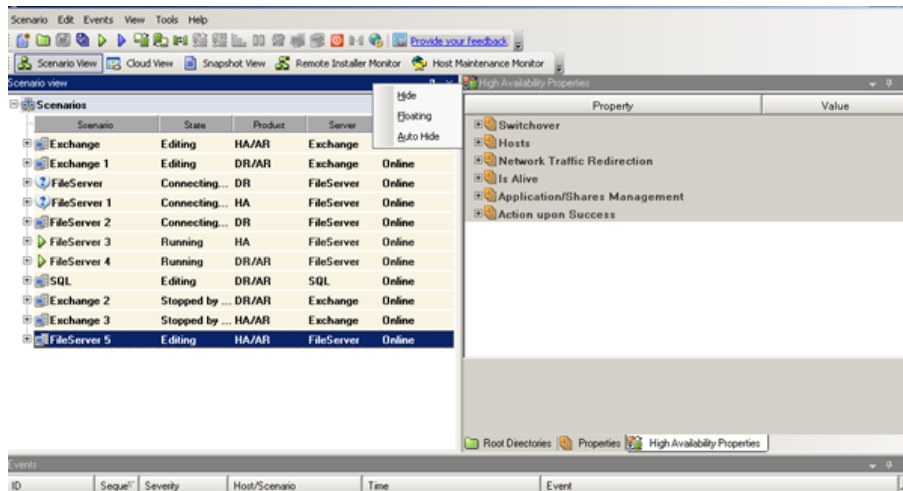


2. If you wish to completely hide the pane, click Hide. If you wish to hide the pane only while you work on other panes, click Auto-Hide. Click its tab to return to the hidden pane.

# Toolbars

Arcserve RHA provides two toolbars to facilitate your work:

- Standard toolbar
- Viewing toolbar

# Standard Toolbar

The Standard toolbar buttons provide quick access to the most commonly used functions in the Manager. The following list contains a brief description of each toolbar option:

**New**

Create a new scenario using the Scenario Creation Wizard. For more information, see Create a New Scenario.

**Group**

Create a new scenario group. For more information, see Create a Scenario Group.

**Save**

Save a selected scenario. For more information, see Save Scenarios.

**Save All**

Save at once all existing scenarios. For more information, see Save Scenarios.

**Run**

Run the selected scenario to start the replication process. For more information, see Initiate Replication.

**Run (Assessment mode)**

Run the selected scenario in Assessment Mode. Refer to How Replication Works.

**Synchronize**

Activate the synchronization process (whether replication is running or not). For more information, see Synchronize Master and Replica Servers.

**Restore Data**

Recover lost or corrupted Master data from any Replica by activating a synchronization process in the reverse direction. For more information, see Recovering Data and Servers.

**Difference Report**

Generate a Difference Report, which displays the difference between a Master and its Replica(s) at a certain point in time. The comparison is performed using the same algorithms that are used in the synchronization process, but no data is transferred. For more information, see Create Difference Reports.

**Perform Switchover**

[For HA only] Switch the active and passive roles between the Master and Replica servers following their current status. For more information, see Switchover.

**Suspend Is Alive Check**

[For HA only] Suspend the Is Alive check that verifies that the active server is operational. For more information, see Is Alive.

**Refresh Statistics**

Update the scenario state information and live statistics display. For more information, see Refresh Statistics Display Manually.

**Suspend Replication**

Suspend replication updates on the Replica host in order to perform system maintenance or some other form of processing that does not modify the replicated data there. Changes continue to be recorded for update on the suspended Replica, but are not actually transferred until replication is resumed. You cannot suspend replication during synchronization. For more information, see Suspend Replication.

**Replica Integrity Testing**

Perform Assured Recovery test in a non-scheduled mode. For more information, see Performing Assured Recovery Test in a Non-Scheduled Mode.

**Start/Stop VM**

Stop or start a virtual machine. This is applicable for a Data Recovery or High Availability full system scenarios.

**Delete All VM Resources**

Delete all temporary VM resources such as disk files, snapshots, and other temporary files.

**Launch Host Maintenance**

Prepare a node in your replicated system to planned maintenance procedures, while avoiding resynchronization once these procedures are completed. For more information, see Host Maintenance.

**Offline Sync Management**

Specify that RHA synchronizes data block by block and then start the replication process.

**Configure HTTP Proxy Server**

Specify proxy server details to connect to RHA engines.

**Provide Your Feedback**

Open the feedback page.

**Help Topics**

Open Online Help.

# Viewing Toolbar

The Viewing toolbar buttons provide quick access to the different windows and monitors in the Manager. The following list contains a brief description of each toolbar option:
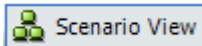


The **Scenario View** gives you access to the main Manager screen, which enables you to create, monitor, and manage replication scenarios.



The **Cloud View** lets you connect to Amazon EC2 to manage your instances.
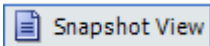


The **Snapshot View** gives you access to the VSS Management window, which enables you to view and manage VSS snapshots.



The **Remote Installer Monitor** gives you access to the Remote Installer view, which enables you to view the installation status of Engines, you installed with the Remote Installer. For more information about the Remote Installer, refer to the *Arcserve RHA Installation Guide*.
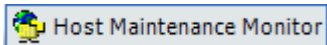


The **Host Maintenance Monitor** gives you access to the Host Maintenance Monitor view, which enables you to view the status of current requests for maintenance preparation.

# Timeout on Web Portal and Manager User Interface

The session timed out without any activity on the Arcserve Replication and High Availability web portal and Manager user interface. As a result, you need to login again. To avoid unexpected time out, you can set timeout value.

**Parameter to set timeout value**

**For Web portal session**

Web portal has Web_UI_Timeout section in INSTALLDIR\Manager\ws_ root\web.config file. The default time is 7200 seconds (2 hours). You can modify the value to adjust this time. After updating the value, restart the control service. The valid range is 1-86400 seconds (1 day).

**For UI session times out**

User interfaces uses the GUITimeout option added in INSTALLDIR\Manager\ mng_core_com.cfg file specifying timeout in minutes. To change this value, remove the # icon placed in front of the attribute, modify the right-hand numeric value, and then restart the control service and GUI. The default value is 120 minutes (2 hours).

# 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

Arcserve RHA protects servers in the context of user-defined scenarios. A scenario is the basic unit of operation and it consists of a definition set that includes:

- The type of application or database server to be protected.
- The type of data protection solution.
- Special tasks, such as Integrity Testing for Assured Recovery.
- The connection details of the Master and Replica hosts.
- The directories, sub-directories, databases and files that will be replicated and their location on the Master and the Replica.
- Configurable properties of the scenario and the Master and Replica hosts, which affect different settings and operations, such as, synchronization method, replication mode, spool size, report and event handling rules, and more.
- Recovery and Switchover/Failover parameters.

Each scenario defines a replication tree that sets the flow of information from the Master server to any number of designated Replicas. It establishes the data recovery procedure, and, if applicable, the switchover parameters. You can configure, add or remove servers from a scenario and select or modify directories. This enables easy, complete control of the replication process over any network, large or small. Each scenario is saved as an XML file. This section describes how to create the following: a scenario group, scenarios using both the Scenario Creation Wizard and a template, and a scenario template.

This section contains the following topics:

# Create a File Server Replication Scenario

The following procedure demonstrates the creation of a generic File Server Replication scenario. For more detailed instructions involving scenarios tailored to specific applications such as Exchange or SQL Servers, see the appropriate *Arcserve RHA Operation Guide*.

**Follow these steps:**

1. Open the Manager. From the Scenario menu, click New or click the New button on the Standard toolbar.

   The Welcome screen of the Scenario Creation Wizard is displayed.
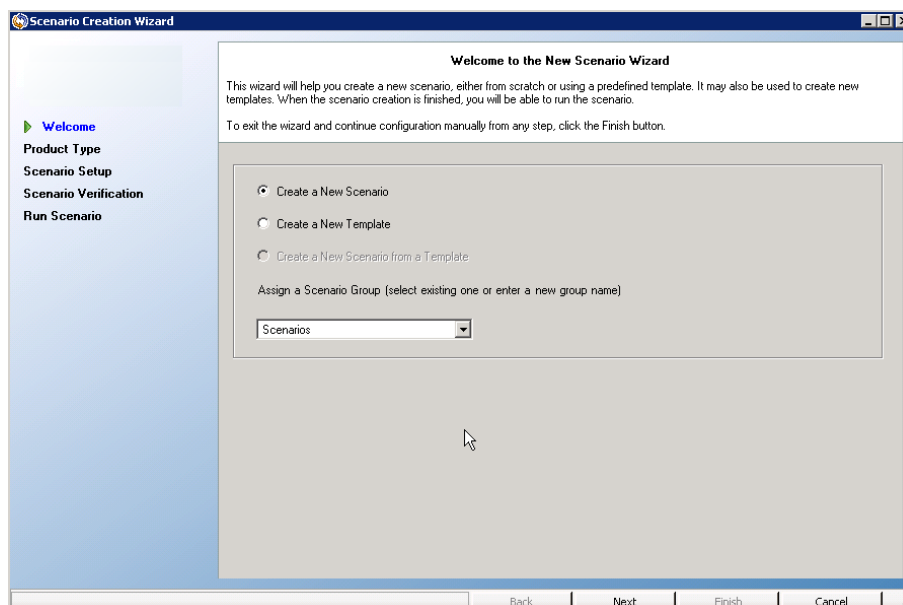


   The Welcome screen enables you to create a scenario, and to assign the new scenario to a scenario group. From the Welcome screen, you can create a scenario directly or from a template, or create a template. For more information about creating templates, see the topic, Using Templates.

   **Note**: To access other Manager features while you are creating a scenario, minimize the Scenario Creation Wizard. The Scenario Creation Wizard is bound to the Scenario View. If you switch views, the wizard is automatically minimized.

2. Select the required options, as follows:

   a. Select Create a New Scenario.

   b. Assign your scenario to the Group, "Scenarios" or type a name to create

a Group.

c. Click Next.

The Select Server and Product Type screen opens.

This screen lists the supported applications and product types available, depending on how you have licensed Arcserve RHA.

3. Select the required Server, Product Type, Integration options, as follows:
   a. From the Select Server Type list, click File Server. For more information about the remaining server types, see the appropriate Operation Guide.
   b. From the Select Product Type list, click Replication and Data Recovery Scenario (DR). For more information about High Availability Scenarios (HA), see the topic, Create a New File Server HA Scenario.
   c. Select Integrity Testing for Assured Recovery if necessary. For more information about Integrity Testing for Assured Recovery, see the topic, File Server Assured Recovery.
   d. From the Integration Options list, select one of the following:

      **None**

      Indicates that you want to create this scenario without integrated Arcserve products. Click Next.

      **Arcserve Backup**

      Indicates that this scenario uses Arcserve Backup to back up the RHA replica.

      Enter the backup server name. Click Next.

      **Note:** For more information about backups, see the Arcserve Backup *Administration Guide* on the Arcserve Backup bookshelf.

The Master and Replica Hosts screen opens.
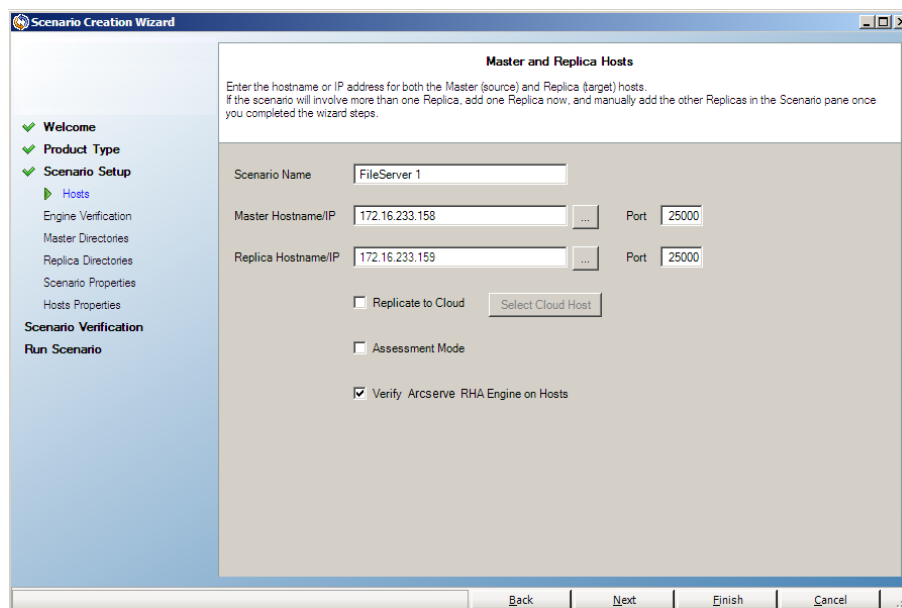
On this screen, specify the host to protect (Master) and the host that holds the replicated data (Replica).

4. Select the required Master and Replica Hosts, as follows:

   a. In the Scenario Name field, accept the default name or enter a unique name.

   b. In the Master Hostname/IP field, enter the hostname or IP address of the Master server. This server is the source server. Use the Browse button to find one.

   c. In the Replica Hostname/IP field, enter the host name or IP address of the Replica server. This server is the target server. Use the Browse button to find one. If you want to include additional Replicas in your scenario, enter the details for the first or most upstream server here. When you have finished the Wizard to create the scenario, you can manually enter additional Replica servers. See the topic, Add Additional Replica Servers.

   d. In the Port fields, accept the default port number (25000) or enter new port numbers for the Master and Replica.

   e. (Optional) If you want to gather statistics about the accurate bandwidth usage and compression ratio benchmarking without actually replicating data, enable the Assessment Mode option. If you select this option, no replication occurs, but a report is provided after the assessment process is completed. For this example, do not enable this option.

f.   (Optional) Enable Verify Arcserve RHA Engine on Hosts to verify whether the Engines are installed and running on the specified Master and Replica hosts. If Engines are not installed on the hosts you specified, you can use this option to install the Engines on one or both hosts remotely. For this example, enable this option.

g.   Click Next.

The Hosts Verification screen opens if you enabled the option, Verify the Arcserve RHA Engine on Hosts. The software verifies the existence and connectivity of the Master and Replica hosts specified on the previous screen. After connections are verified, the software checks whether an Engine is installed on each host. If you log in to the Manager with different user credentials than remote hosts, the Server Status is reported as Not Connected. You are then prompted to enter User Credentials for each selected host. Verification repeats after you do so.
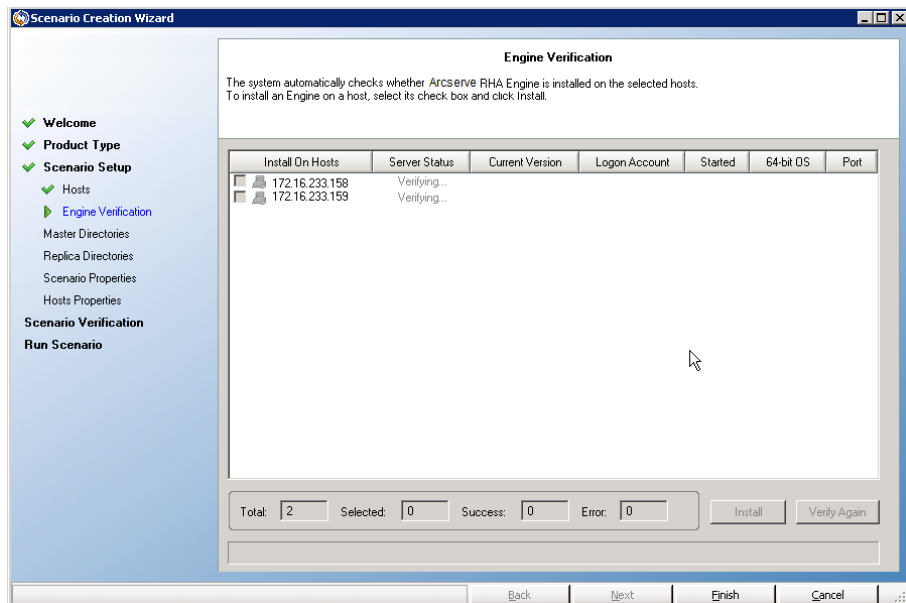


5.   From the Hosts Verification screen, check whether an Engine is installed on the selected hosts using the Current Version column.

Do one of the following:

‒   If an Installed indication appears under the Server Status column in both rows, you can move to the next page.

‒   If an Installed indication appears, but the version is different from the version of the Control Service you are using, install the current version.

–  If a Not Installed indication appears, install the Engine. Click Install to install the Engine on the selected host. You can install the Engine on both hosts at once. Click the servers and click Install.

After you click Install, you are prompted to enter the Arcserve RHA Engine service account credentials:

- For Replication scenarios, it is sufficient to be a Local Administrator (Local System).

- For Clusters (Including replication scenarios), run under the same account as the Cluster Service account.

- For HA scenarios, run under an account with Domain Admin-istrative privileges in Active Directory environment or have local administrative privileges in workgroup environment.

a. Wait while installation completes and the Engine version number appears in the Current Version column.

b. Click Next.

The Master Root Directories screen opens.

This screen displays the directories and files on the Master server. These directories and files are the data that can be replicated and protected. The software automatically aggregates data with a common path into one directory.

When you select root directories for the Master and Replica, the character length of the root directory and subdirectory names cannot not exceed 1024 bytes.

6. From the Master Root Directories screen, select the directories and files you want to replicate from the Master to the Replica by clicking their check boxes. You can exclude folders and files by clearing check boxes.

When you select a drive or directory on the left, the software displays its contents on the right. For more information, see Filter Master Directory Files.

Replication of mount points succeeds only if those files were added to the Master before the Engine began to run. If you included the mount points in the Master root directories when the Engine was running, no error is reported but replication does not start. In this case, restart the Engine to initiate replication.

7. Click Next. The Replica Root Directories screen opens.

Accept the default or type a new directory name.

8. Click Next. The Scenario Properties screen opens.
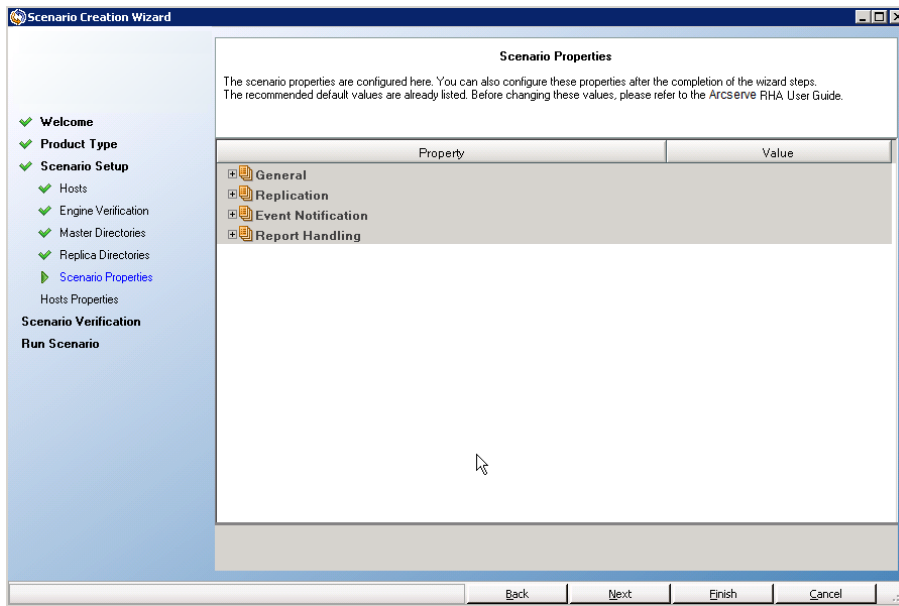
9. From the Scenario Properties screen, configure the properties that affect the entire scenario. For this example, simply accept the defaults. These properties can also be configured outside of the Wizard. For more information about configuring scenario properties, see the topic, Configure Scenario Properties. Click Next. The Master and Replica Properties screen opens.



10. From the Master and Replica Properties screen, configure the properties that are related to either the Master or Replica hosts. For this example, simply accept the defaults. For more information, see the topic, Configure Master or Replica Server Properties.

Review Spool Information before you change any Spool properties.

Click Next. Wait until the Scenario Verification screen opens.

11. The software validates the new scenario and verifies the parameters for a successful replication. Once verification completes, the screen opens, displaying any problems and warnings. The software permits you to continue even if warnings are displayed. Resolve any warnings to for a proper software operation.

    Click Next when all errors and warnings are resolved. The Scenario Run screen opens.

12. Running the scenario initiates the data synchronization process. Select Run Now to start synchronization immediately or Finish, which saves the scenario configuration and allows you to initiate synchronization later.

    **Note**: Synchronization takes a while, depending upon your data size and network bandwidth. Select Offline Synchronization if necessary. For more information, see the topic, Synchronization Methods.

    If you select Run Now, the software notifies you when synchronization completes. Now the real-time replication is operational and the replication scenario is active.

    A synchronization report is generated. To view the report, see the topic, Viewing a Report.

**Considerations for Creating the Arcserve D2D Scenario**

When you select the Arcserve D2D as the scenario type, consider the following limitations and recommendations:

- When the D2D destination changes and you update the Arcserve RHA scenario with the new destination, old sessions on Replica are lost after synchronization. The sessions are lost because those old sessions do not exist on the new D2D destination.

- When the D2D destination is a local disk, verify that the Arcserve RHA spool is configured to a volume which is not protected by Arcserve D2D. Otherwise Arcserve D2D will backup the journal files stored in the Arcserve RHA spool.

# How to Create Arcserve Central Host-Based VM Backup Scenarios

The following diagram illustrates an installation setup having RHA, D2D, and Central Host-Based VM Backup.
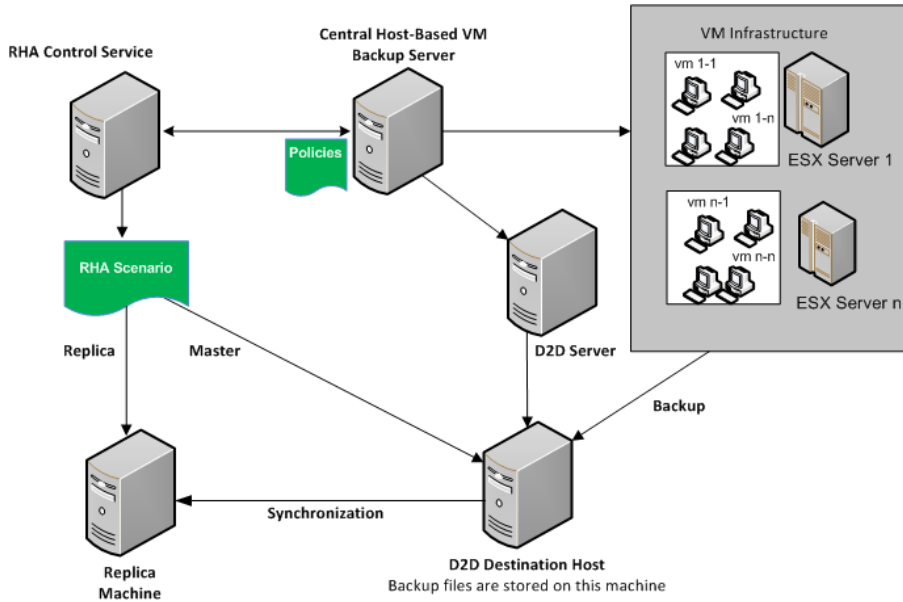


To protect multiple backup files (these files are created by Arcserve D2D), you create a Central Host-Based VM Backup scenario. Specify the Arcserve Host-Based VM Backup server credentials and select virtual machines from the D2D host server.

The following diagram illustrates the process to create a Arcserve Central Host-Based Backup scenario.
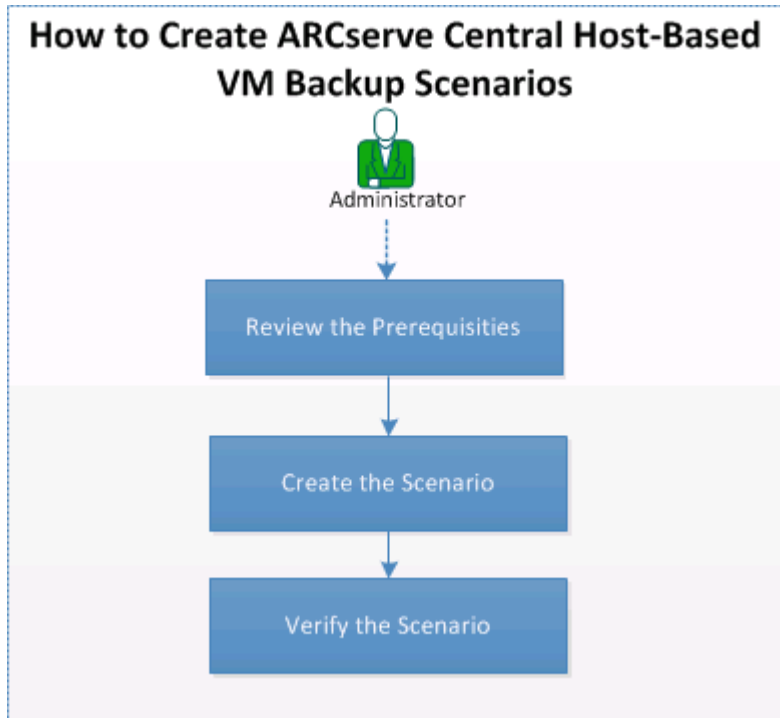
Perform the following tasks to create a Arcserve Central Host-Based Backup scenario:

1. Review the Prerequisites

2. Create the Scenario

3. Verify the Scenario

## Review the Prerequisites

Make sure that you have reviewed the following prerequisites before you create the Arcserve Central Host-Based VM Backup scenario:

- You have Arcserve D2D installed and the backup environment is set up. For information about Arcserve D2D, refer to the *Arcserve D2D User Guide.*

- You have Arcserve Central Host-Based VM Backup (HBBU) set up and the Central Host-Based VM Backup server can access D2D servers. For information about Arcserve Central Host-Based VM Backup, refer to the *Arcserve Central Host-Based VM Backup User Guide.*

- You know how to create RHA file server scenarios.

# Create the Scenario

You create the Central Host-Based VM Backup Integration scenario when you want to replicate the Arcserve created backup files.

**Follow these steps:**

1. Open the Manager. From the Scenario menu, click New or click the New button on the Standard toolbar.

   The Welcome screen of the Scenario Creation Wizard is displayed.

2. Select Create a New Scenario.

   The Select Server and Product Type screen opens.
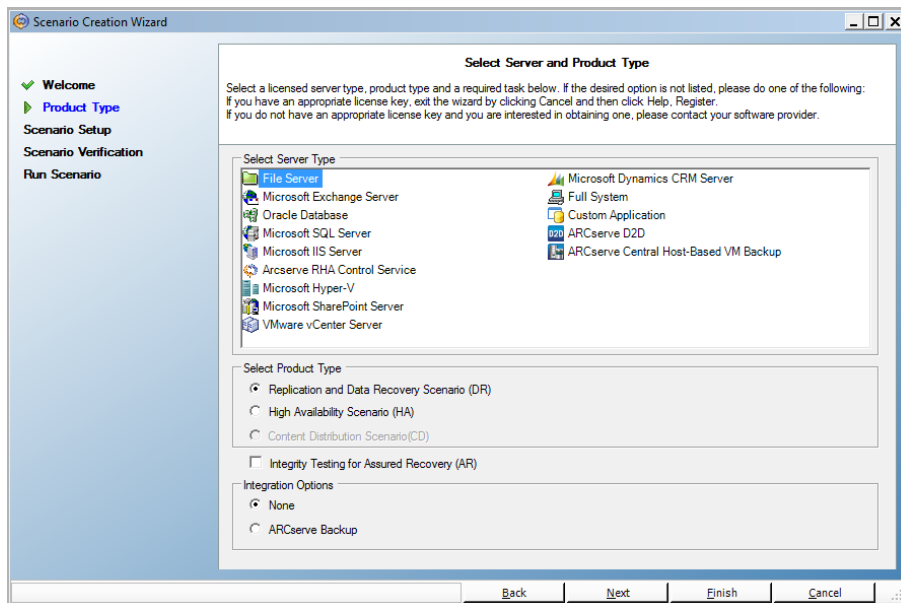


3. Select the following options and click Next.

   a. Server Type: Arcserve Central Host-Based VM Backup.

   b. Product Type: Replication and Data Recovery Scenario (DR).

   c. Integration Options: None.

   The Arcserve Central Host-Based Backup server credentials dialog opens.
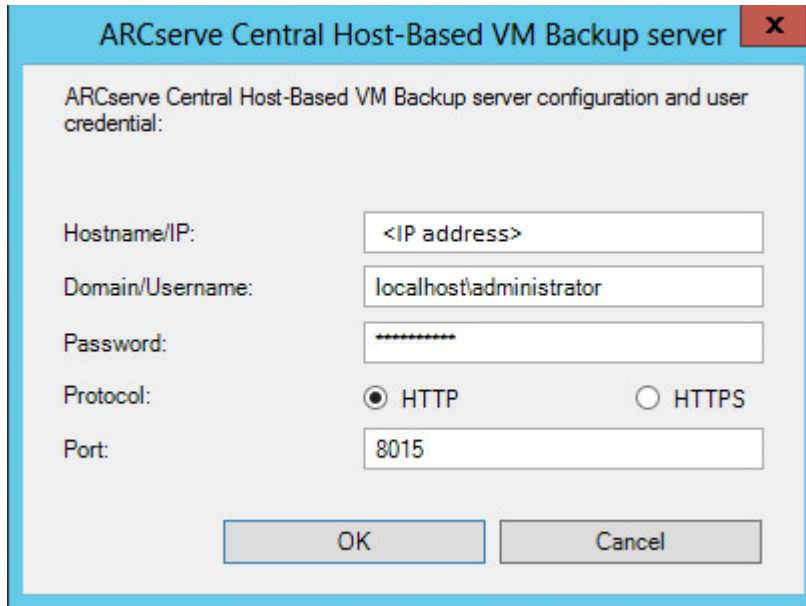
4. Enter the Central Host server credentials and click OK. The server name is populated based on what you enter in step 3.

The Arcserve Central Host-Based VM Backup Destination Host and Virtual Machine Selection screen opens.
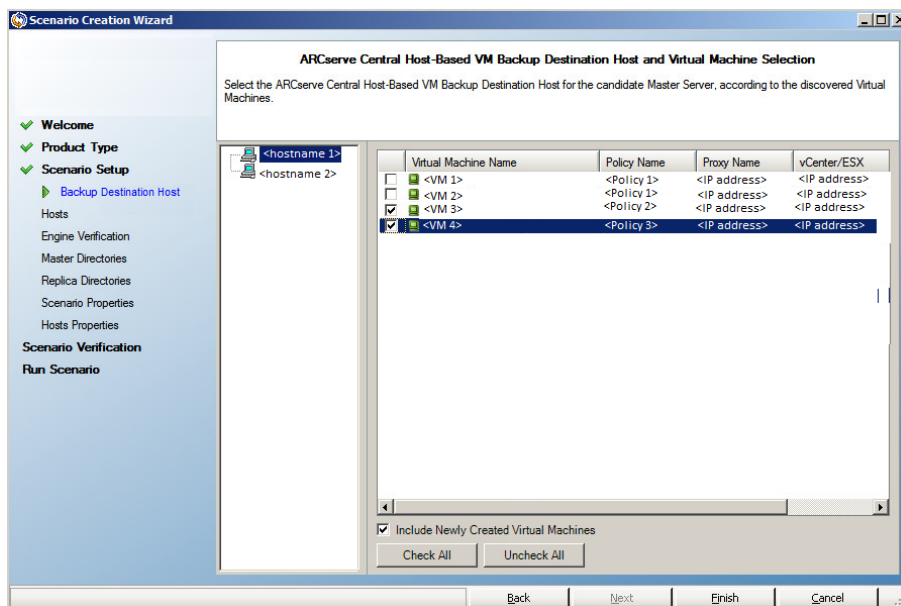


Arcserve RHA connects to the Central Host-Based Backup server to get the policy and display backup destination hosts and its virtual machines.

5. Select the hostname and select virtual machines that you want to protect.

**Include Newly Created Virtual Machines:** Specifies that all sub folders in the main host backup folder are replicated when you run this scenario. Any newly

created VM backup folders are also replicated. Only the folders for unchecked VMs are excluded. Those folders are marked as excluded folders. If you do not select this option, only the selected backup folders are replicated.

The backup files of the selected virtual machines are replicated when you run this scenario. These are the backup files that are created by Arcserve D2D.

6.  Enter the following Master and Replica details:

    **Scenario Name:** Accept the default name or enter a unique name.

    **Master Hostname/IP:** Automatically populated based your hostname selection.

    **Replica Hostname/IP:** Enter the host name or IP address of the Replica server. This server is the target server. Use the Browse button to find one.

    **Port:** Accept the default port number (25000) or enter new port numbers for the Master and Replica.

    **(Optional) Verify Arcserve RHA Engine on Hosts:** Enable to verify whether the Engines are installed and running on the specified Master and Replica hosts.
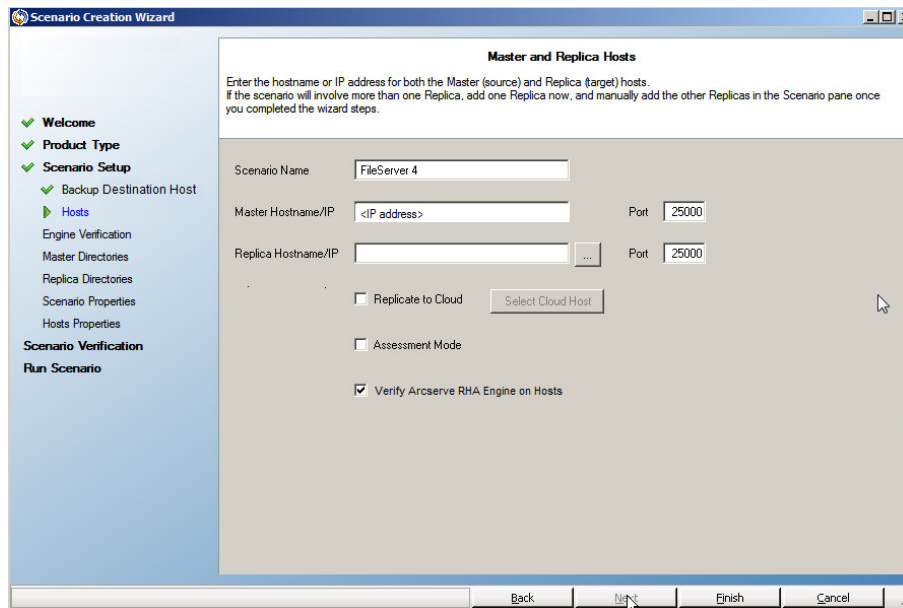


7.  Click Next.

    The Engine Verification screen opens.

    The Hosts Verification screen opens if you enabled the option, Verify the Arcserve RHA Engine on Hosts. The software verifies the existence and connectivity of the Master and Replica hosts specified on the previous screen.

8.  Click Next.

    The Master Root Directories screen opens.
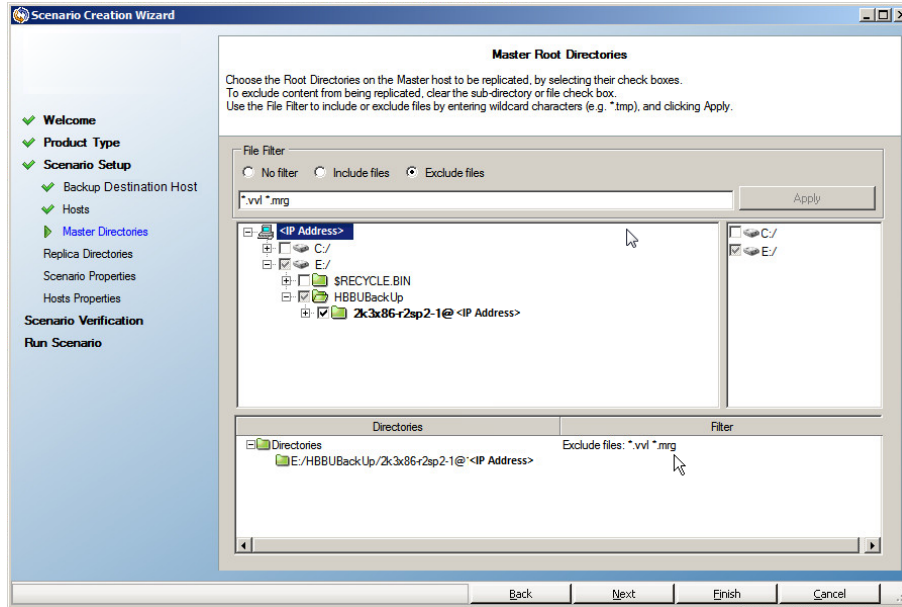
The RHA engine discovers the backup folders of the selected virtual machines. These backup folders are automatically selected.

**Note:** These folders are the backup folders that are created by Arcserve D2D.

When you select Include Newly Created Virtual Machines in the"Arcserve Central Host-Based VM Backup Destination and Virtual Machine Selection" screen, the main backup folder is selected for replication and the excluded folders are listed in the filter pane.
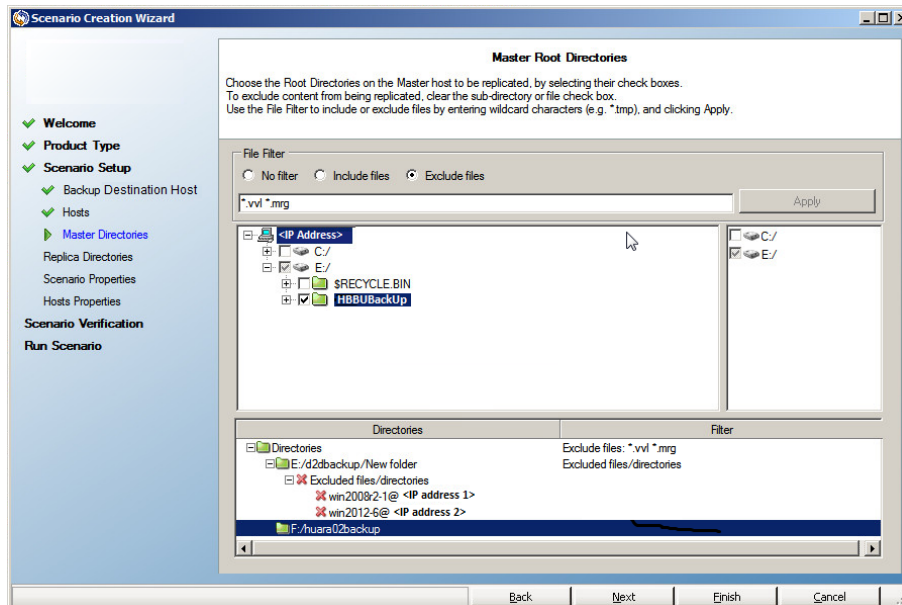


9.  Click Next.

The Replica Root Directories screen opens.

10. Accept the default and click Next.

    The Scenario Properties screen opens.

11. Configure the properties that affect the entire scenario. For this example, simply accept the defaults. These properties can also be configured outside of the Wizard. For more information about configuring scenario properties, see Configure Scenario Properties.
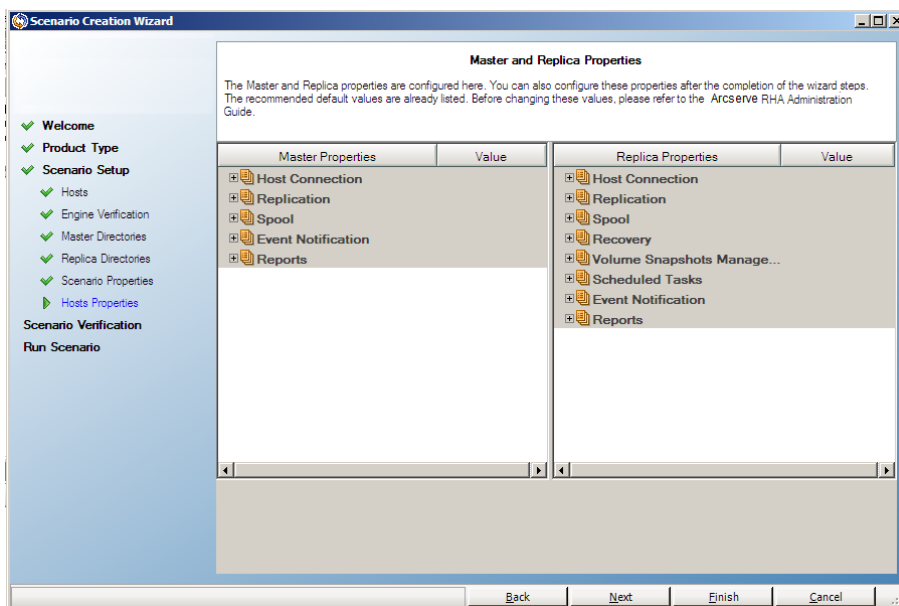
12. Click Next.

    The Master and Replica Properties screen opens.
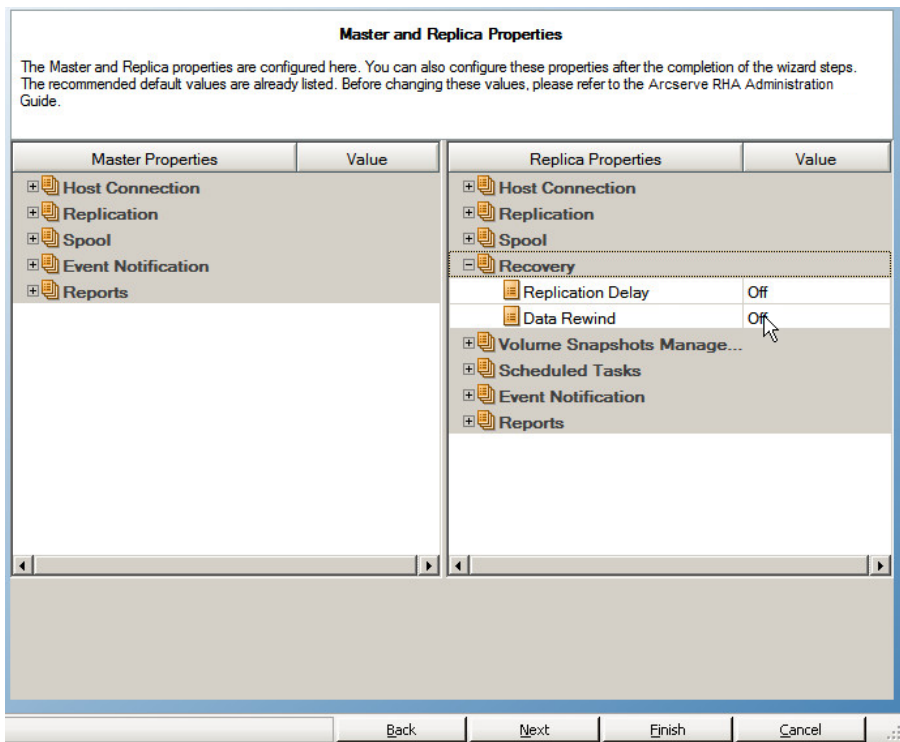


13. Configure the properties that are related to either the Master or Replica hosts. For this example, simply accept the defaults. For more information on how to configure master and replica properties, see Configure Master or Replica Server Properties.

14. Click Next.

    The Scenario Verification screen opens.

    The software validates the new scenario and verifies the parameters for a successful replication. After verification completes, the screen opens, displaying any problems and warnings. The software permits you to continue even if warnings are displayed. Resolve any warnings if required.

15. Click Next when all errors and warnings are resolved.

The Scenario Run screen opens.

16.  Select Finish.

The Arcserve serve Host-Based VM Backup scenario is successfully created. Now you can run this scenario and back up virtual machine files that are created by Arcserve D2D.

## Verify the Scenario

After the Arcserve Central Host-Based VM Backup scenario is created, the scenario is listed on the RHA Manager, Scenario View. Verify that the scenario is listed in the Scenario view with correct details.

# Create a New File Server HA Scenario

Before you start this procedure, read the IP Redirection section and perform the prerequisite steps necessary to protect your environment. If you decide to use the Move IP redirection method, you must add a new IP Address to the Master NIC before you create scenarios.

This procedure launches a Wizard that guides you through the steps required for HA scenario creation. However, properties can also be configured outside of the wizard.

**To create a new file server HA Scenario**

1. Open Manager and choose Scenario, New or click the New Scenario button to launch the wizard.

   The Welcome dialog opens.

2. Choose Create a New Scenario, select a Group from the list and then click Next.

   The Select Server and Product Type dialog opens.
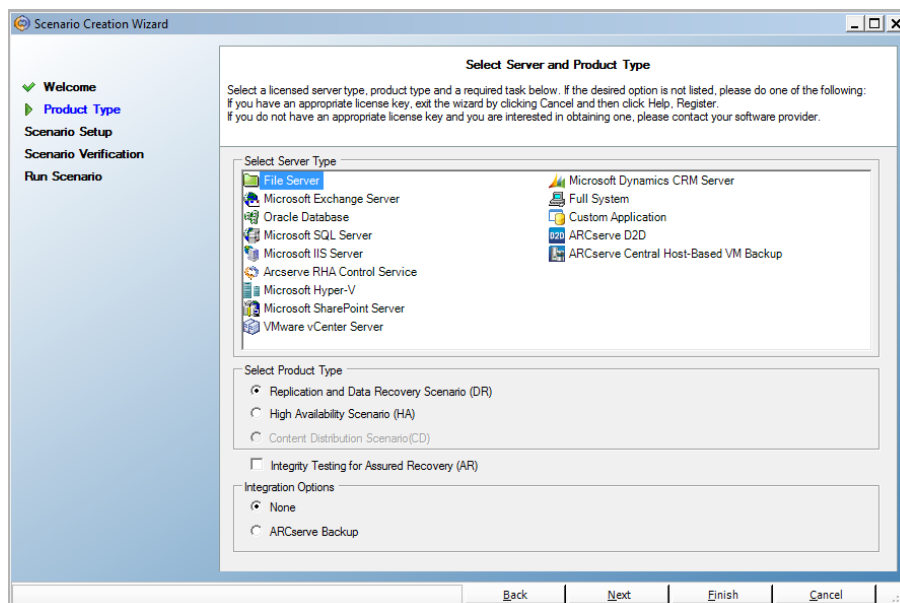


3. Choose File Server, High Availability Scenario (HA) and then click Next.

   The Master and Replica Hosts dialog opens.
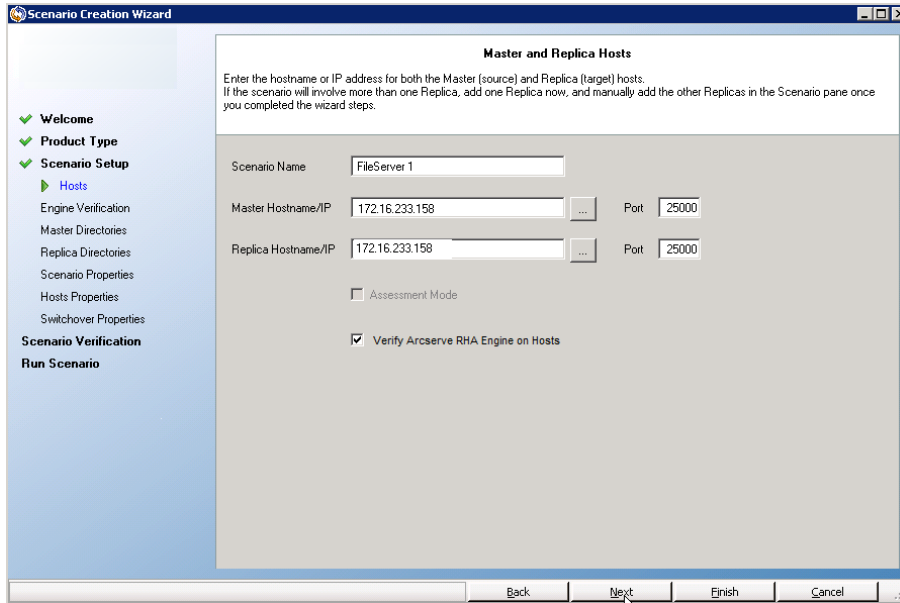
4. Type a Scenario Name, enter the Hostname or IP Address and Port number for both Master and Replica servers, make sure the Verify Arcserve RHA Engine on Hosts option is enabled (default) and click Next.
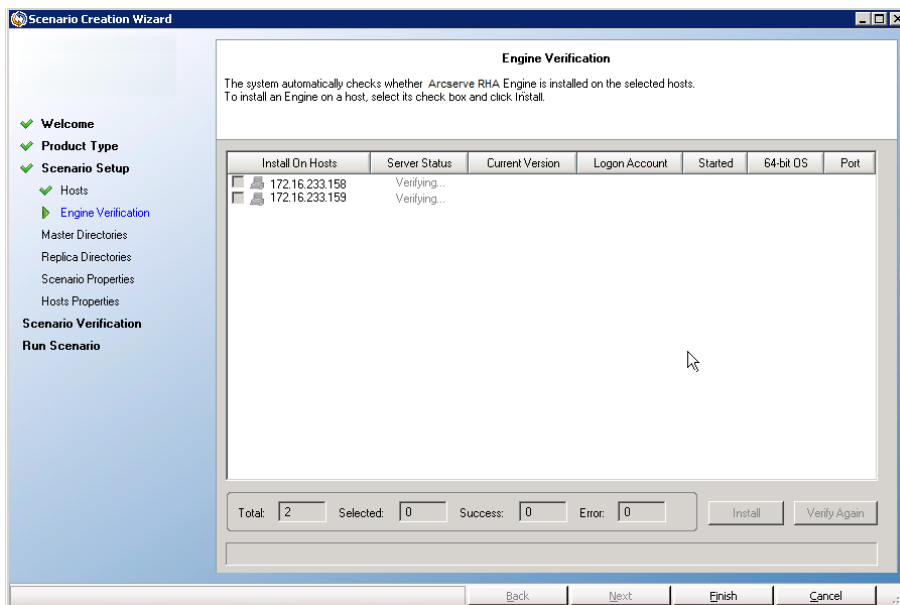


5. Wait for Engine Verification to complete and click Next. If required, click Install to upgrade the Engine on one or both servers and Verify Again.

   The Master Root Directories dialog opens, with a list of auto-discovered directories. By default, nothing is selected. Expand folders and select the data you wish to protect. Note that not all system files can be selected and are excluded by default.

6. Click Next. The Replica Root Directories screen opens.

Accept the default root directory or type a new name.

7.  Click Next. The Scenario Properties dialog opens.
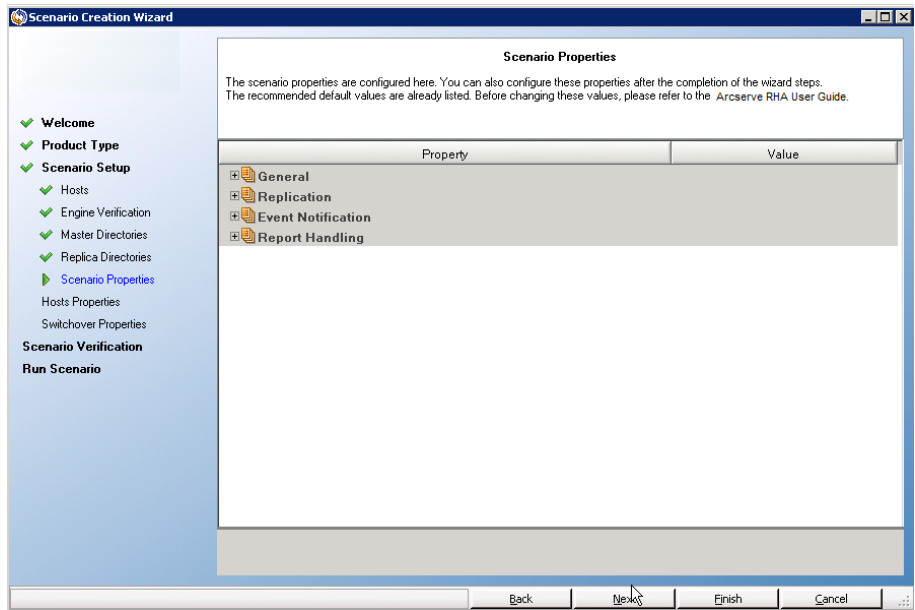


8.  Scenario properties control the entire scenario. Accept the default values or set new values as required. Click Next when done. These properties can also be configured outside of the Wizard. For more information, see the topic, Configuring Scenario Properties.
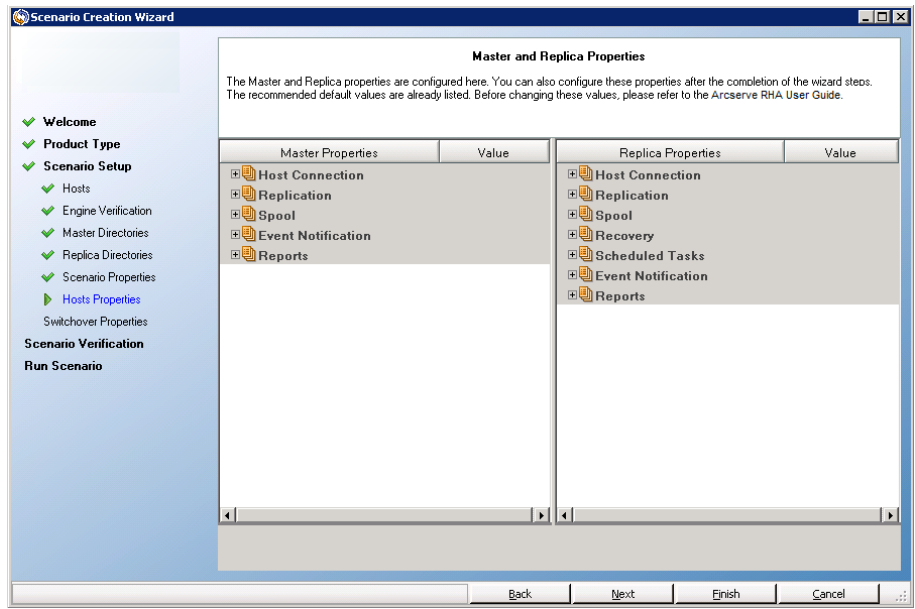
    The Master and Replica Properties dialog opens.



9.  Master and Replica properties apply to only host servers. Accept the default values or change values, as desired. Click Next.

    Wait for the Switchover Properties dialog to retrieve information.

10. Set the desired Network Traffic Redirection method, as described in the IP Redirection section. Click Next.

The Switchover and Reverse Replication Initiation dialog opens.
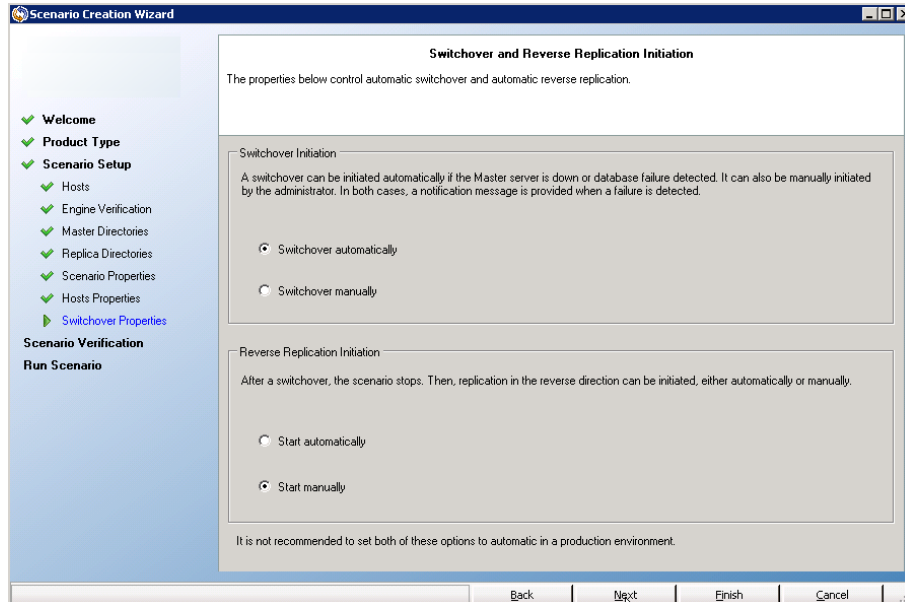


11. Choose the desired start options. For File Server scenarios, set reverse replication initiation to manual. Automatic is not recommended. For more information, see Switchover Considerations. Click Next.

Wait for Scenario Verification to complete.

12. If errors or warnings are listed, resolve them before continuing. When ready, click Next.

The Scenario Run dialog opens.

13. Click Run Now to start synchronization and activate the scenario, or click Finish to run the scenario later.

# Use Scenario Groups

Each scenario is assigned to a default scenario group called **Scenarios**. You can use this group for all scenarios you create, or you can add new groups to organize your scenarios according to your own criteria. Scenario groups are displayed on both the Manager and the Overview Page.

In distributed server environments, in which several servers (database server, application server, web front end server) comprise the environment, you must create individual scenarios to protect all the servers in the deployment. If an Is Alive check triggers a failover, only the affected server is failed over to its Replica. The resulting data split, where some operations are applied to original Master servers and other operations are applied to the Replica in the failed scenario, can introduce performance issues.

Scenario groups let you manage related scenarios, such as those protecting all servers in a distributed environment, as a single entity. For example, for end-to-end protection in a distributed server environment, you may have a SQL scenario to protect the database component, and several application-specific scenarios to protect application servers. A scenario group lets you set switchover properties at the group level, instead of at individual server levels.

For more information, see the topic, Enable Scenario Group Management, and the Operation Guide for the specific distributed server application.

**Note**: For SharePoint Server farms, scenario group creation is handled automatically. For other distributed server environments (BlackBerry Enterprise Server, Microsoft Dynamics CRM), you must manually create groups and scenarios.

**Next Steps:**

- Create a Scenario Group
- Enable Scenario Group Management
- Run a Scenario Group
- Stop a Scenario Group

# Create a Scenario Group

There are two ways to create a scenario group:

- During the creation of a new scenario, through the Scenario Creation Wizard.

- Before the scenario creation, through the **New Group** option, as described below.

**Note**: We recommend planning and creating the scenario groups you want to use in advance. After you assign a scenario to a certain group, you cannot move it to another group.

**To create a new scenario group**

1. From the Manager, click Scenario, New Group from the menu, or click the

   New group [icon] button on the Standard toolbar.

   A New group folder is added to the Scenario pane.

2. You can change the group's name by right-clicking it and selecting Rename from the pop-up menu, or by double-clicking the current name and entering a new name.

   The new group name appears on the following places: the Scenario pane, the Group drop-down list in the Scenario Creation Wizard, and the Overview Page.

   **Note**: When no scenario is defined, empty scenario groups do not appear on the Overview Page.

# Set Group Properties

Configure the group properties on the Arcserve RHA Manager Properties tab.

The group properties include:

**Scenario Dependencies**

Manages the interdependencies between scenarios. Usually a distributed application has multiple components/roles/servers which are interdependent. Any scenario can be configured to depend on one or more scenarios or multiple scenarios can depend on a single scenario. These services can be handled by the Scenario Dependencies property.

**Switchover Settings**

Manages the switchover setting in a distributed group. Some of the switchover setting options include:

- Switchover as a Group: If this option is set to On, the whole group (all the scenarios) will be automatically switched over together in case one of the scenarios is failed and ready to take switchover.

- Failure Triggers Group Switchover: A single failure can trigger a group switchover. By default, all scenarios can trigger group switchover, and you can configure some light-weight scenarios to be set to Off.

- Execute Scenario Switchover Settings: This option decides whether the scenario should execute its own switchover settings.

**Scenario Availability Sets**

A distributed application might configure two or more servers to provide same services to improve availability or performance. When one server is down, the other servers are still working and can still provide the services. Scenario Availability Set will be used when Arcserve RHA manages these servers/scenarios of that distributed application.

If two scenarios are configured in the same Scenario Availability Set, the group switchover is initiated only when both scenarios fail. This function is not invoked when one of the options fail.

**Note:** The same group can have one or more Scenario Availability Sets, but one scenario cannot be configured in two different sets.

# Enable Scenario Group Management

Scenario Group Management lets you manage related HA scenarios as a single entity. Switchover can be configured so that when one server fails, all servers in the scenario group switch over at once, alleviating the data split issue. Scenario Group Management applies only to high availability scenarios.

**Note**: For SharePoint Server Farm scenarios, group creation and central scenario management are automatically enabled during scenario creation. For more information, see the *SharePoint Server Operation Guide.* For all other distributed environments, you must manually create the required scenarios, assign each to the same group, and enable group management.

**To enable Scenario Group Management**

1.  From the Manager, right-click the name of the scenario group you wish to manage centrally.

2.  From the short-cut menu, click Enable Scenario Group Management.
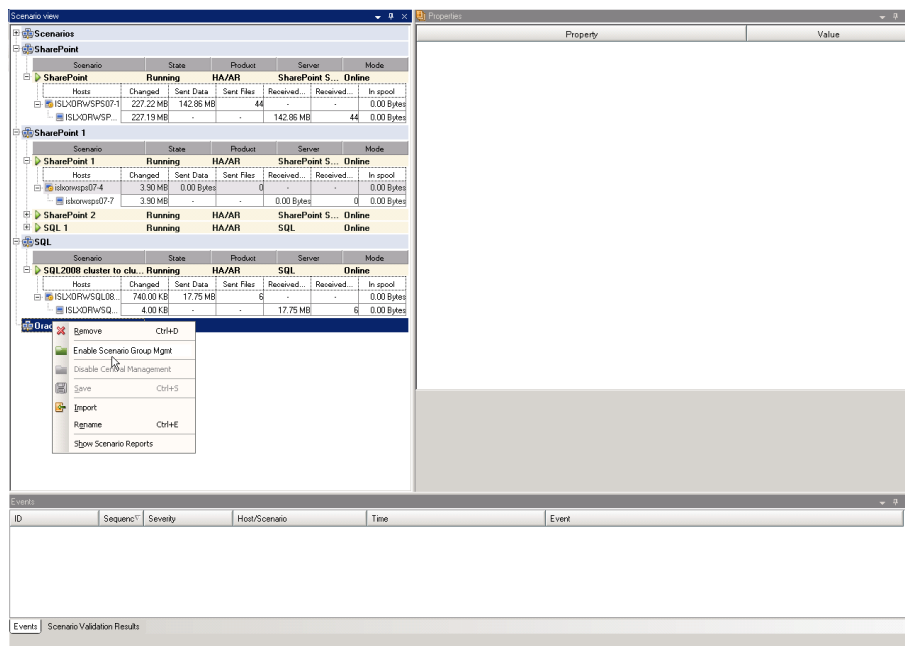


A confirmation message appears. Click OK to continue.

3.  Click the Properties tab and expand the Switchover Settings properties group.

4.  If you want to switch over the whole group as a single entity, set the Switchover as a Group property to On.

5.  Expand the Failure Triggers Group Switchover property and set the value to On for each scenario in the group that should trigger switchover when it fails.
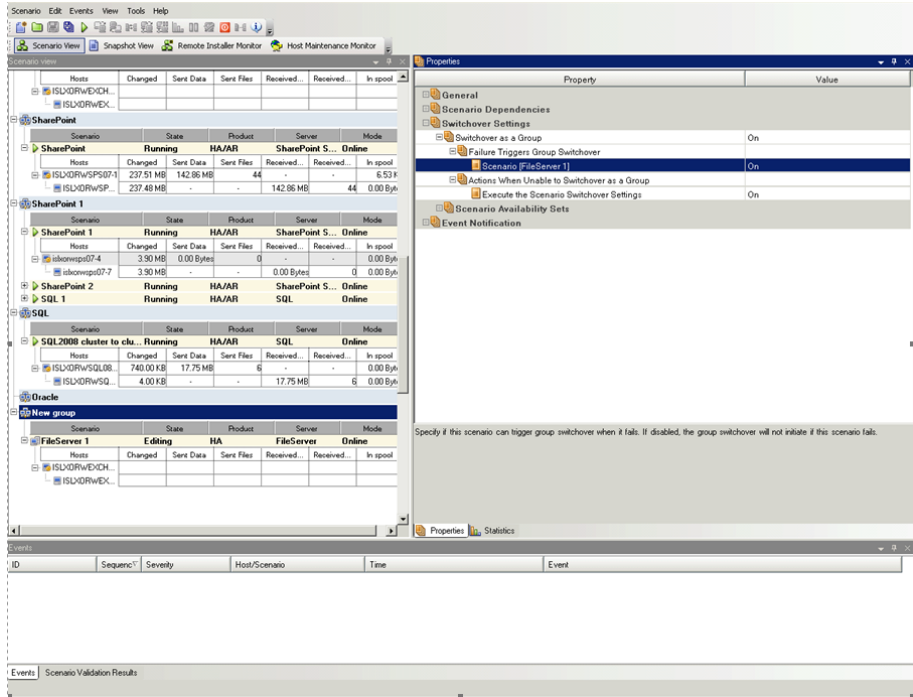
6. Expand the Scenario Availability Sets property. If all servers listed in this property fail, the entire group switches over. Add the names of the scenario group that you wish to monitor and then select the scenarios in that group that will trigger group switchover.
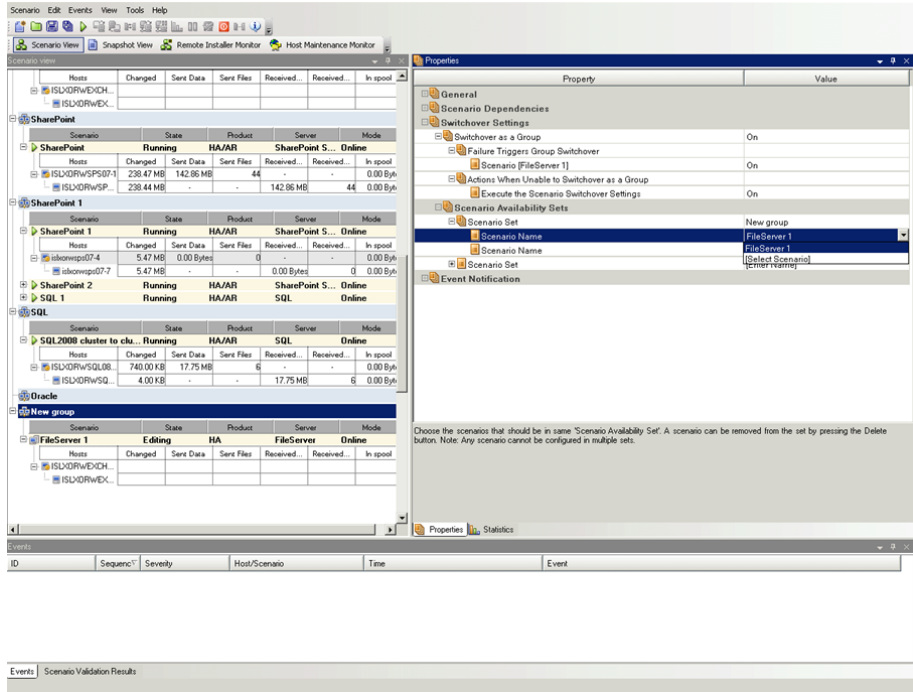
# Run a Scenario Group

Before you can run a scenario group, Arcserve RHA performs a pre-run verification of each scenario in the group and reports any errors or warnings. Each scenario in the group must pass verification for the group to run.

For more information, see [Running the Replication Process](#).

**To run a scenario group**

1. When pre-run verification is successful, click Run Now to run the whole group.

   The Run dialog opens.

2. Choose a synchronization method and click OK. By default, the synchronization method for the group is set to use the method selected for each individual scenario within it. Or, you may apply a method to all scenarios.

   The status for all scenarios in the group changes to Running.

# Stop a Scenario Group

You must stop a group that is currently running if you want to add or remove scenarios. To stop a group, you must stop all scenarios in that group. Click Stop on the Manager toolbar for each scenario in sequence. There is no failure recorded by stopping scenarios.

For more information, see Stop Replication.

# How to Use Templates

Templates are a powerful facility for customizing Arcserve RHA for your own environment. Many applications allow the default values of individual parameters to be modified. Examples might be the default font to be used in new documents in Microsoft Word, or the default signature for new messages in a mail client application. Templates take this idea one step further.

Rather than provide a method to modify a single, global default value, templates offer the ability to create an entire scenario that can be used as the starting point for new scenarios in the future. These special template scenarios contain all the parameters of an actual scenario and all of them may be modified except those that clearly apply to a specific, individual scenario (such as the host names of the Master and Replica servers).

The second important advantage of templates over a set of global defaults is that they allow different sets of defaults to be created for different types of scenarios. For example, the default values appropriate for your Exchange Server High Availability scenarios are not identical to those for your File Server Replication scenarios. With templates, you can create default settings and maintain them individually for each type of scenario your IT environment needs.

This section contains the following topics:

- Create a New Template
- Create a New Scenario using an Existing Template

# Create a New Template

Templates are simple to create and use, and their creation process is basically similar to the creation of a new scenario. However, because a template is not associated with any actual servers, some values cannot be entered, such as the host names or IP addresses of the Master and Replica servers. Also, while default folder paths can be entered on the Directories tab, they must be typed explicitly rather than entered through a file browser.

All the templates are automatically stored in the **Templates** folder on the Scenario pane. This folder does not appear in the Scenario pane until at least one template is created.

**To create a new template**

1. Open the Scenario Creation Wizard by clicking the **New** button on the Standard toolbar, or selecting **New** from the **Scenario** menu.

    The **Scenario Creation Wizard** opens.
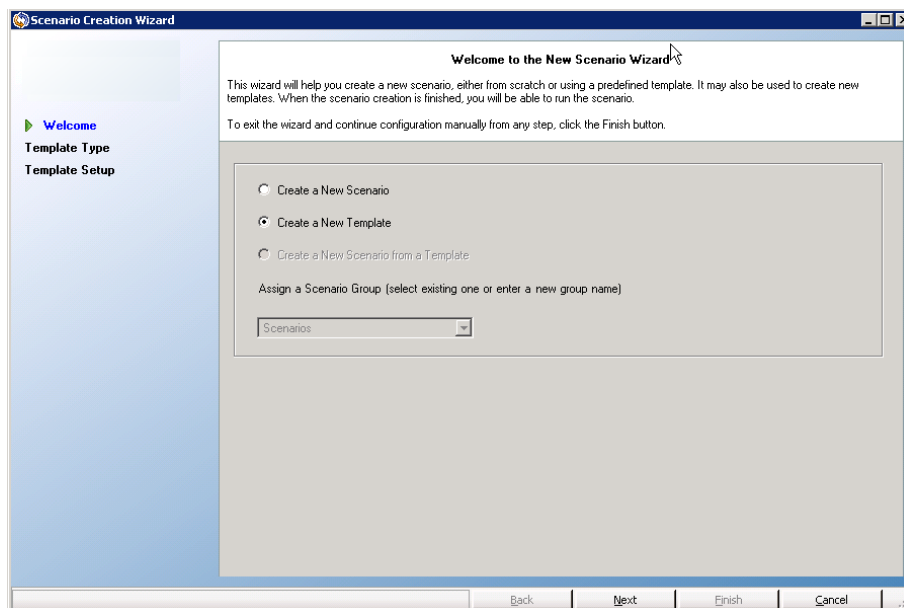


2. Select the **Create a New Template** option button, and click **Next**.

    The **Select Server and Product Type** page opens.
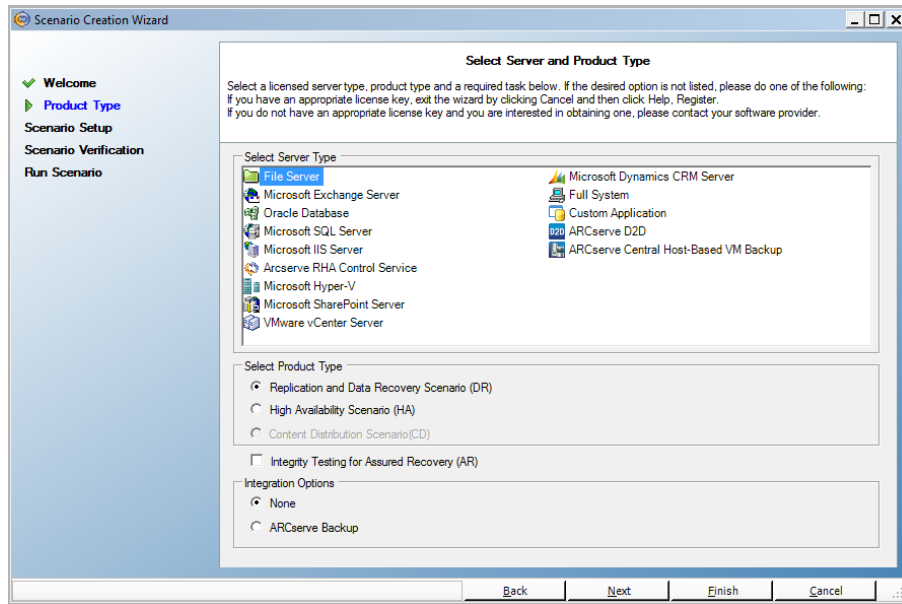
3. Select the required scenario options, and click **Next**.

   The **Template Name and Comments** page opens.
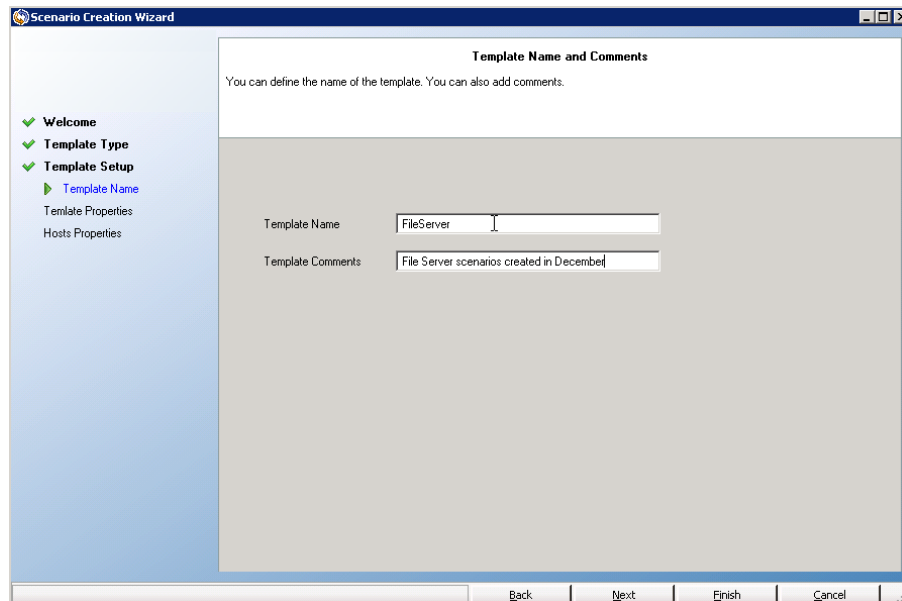


4. Provide a name and description for the template.

5. Click **Next.** The **Template Properties** page opens.
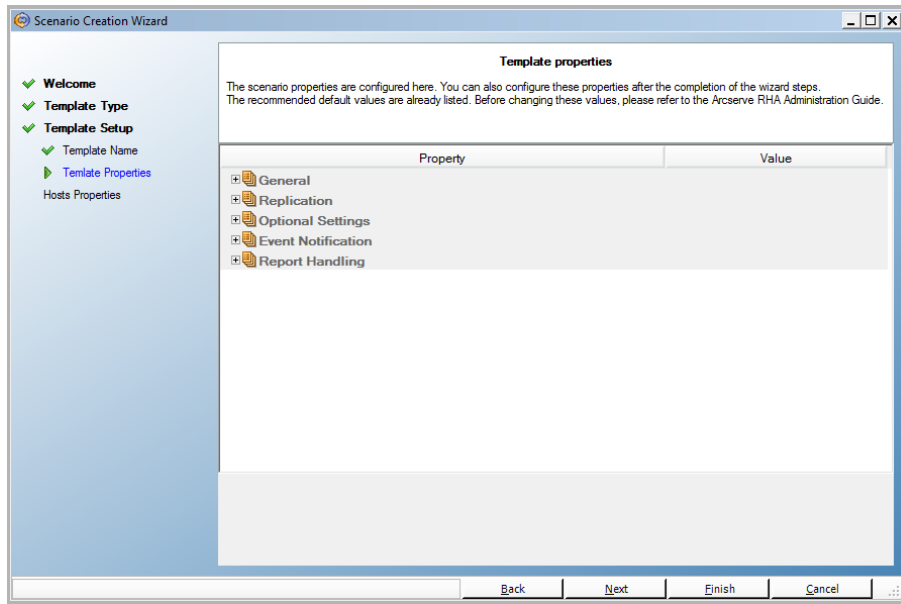
6.  From this step, the wizard pages are similar to the ones you use in creating a new scenario. Continue defining the template by following the wizard instructions.

# Create a New Scenario using an Existing Template

Creating a new scenario using a template saves you the need to separately configure each required scenario. When you are using one of the existing templates, a new scenario is created with all parameter values taken from that template.

**Note:** Scenarios cannot be created from a template until the template has been saved. Changing parameter values in a template will not change the values of those parameters in a scenario that was previously created from it.

**To create a new scenario using an existing template**

1. Open the Scenario Creation Wizard by clicking the **New** button on the Standard toolbar, or selecting **New** from the **Scenario** menu.

   The **Scenario Creation Wizard** opens.

2. Select the **Create a New Scenario from a Template** option button, and assign the new scenario to a group. Then, click **Next**.

   The **Select Template** page opens displaying a list of available templates.
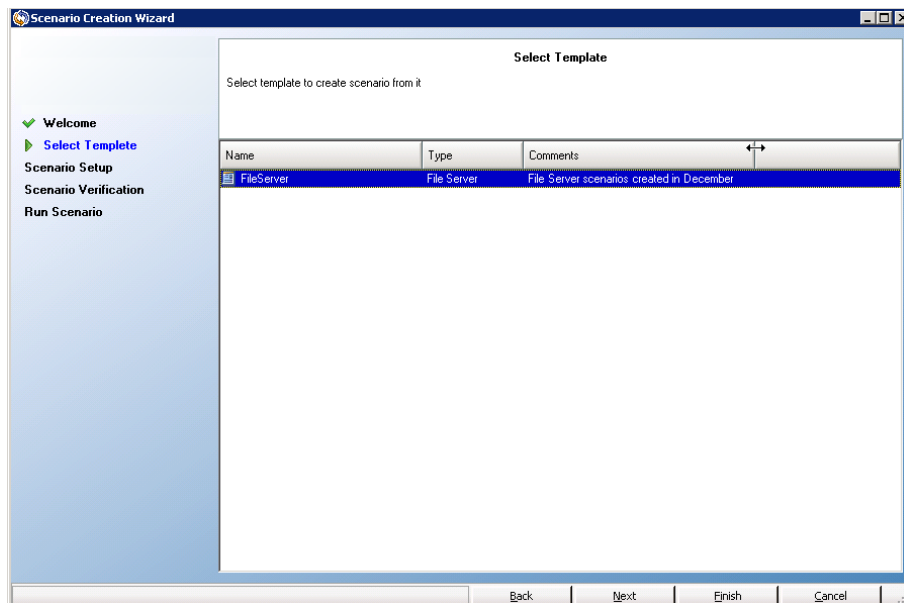


   The available templates appear in this page.

3. Select the template that is best suited for your needs and click **Next**.

   The **Master and Replica Hosts** page opens.

4. The default scenario name is the template name. You can either keep it or change it.

From this step, the wizard pages are similar to the ones you use in creating a new scenario. Continue defining the new scenario by following the wizard instructions specific to the application you are protecting. For more information, see the appropriate Operation Guide.

# Managing Hosts that use a NAT Device

The RHA Control Service is responsible for the management of all scenario-related tasks. The Control Service communicates with the RHA Engines on all hosts participating in a scenario. When your network setup is such that the Master or Replica or both use a Network Address Translation (NAT) device to connect to a public domain. In that case, the Control Service cannot communicate with those servers. To enable RHA Control Service to communicate with such servers, use the Arcserve RHA NAT utility. You configure the NAT utility by providing the IP address and port details of the Control Service and servers in the public domain.

The following diagram illustrates how the RHA Control Service communicates with servers behind a NAT device.



Creating a Scenario for a Host behind a NAT Device
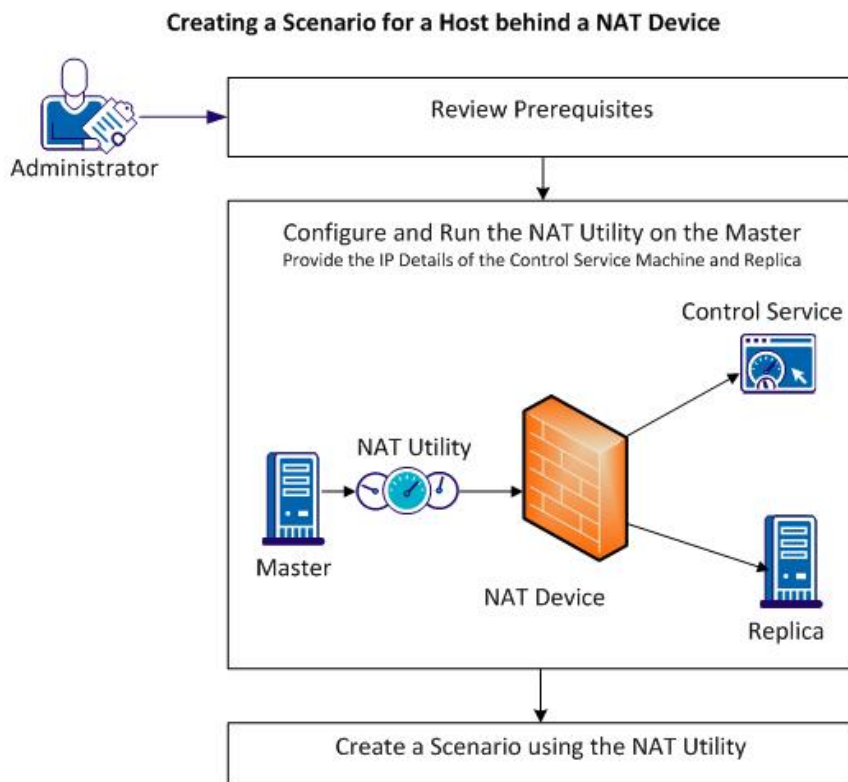
Perform the following tasks to use the NAT utility:

- Configure the NAT Utility
- Create a Scenario using the NAT Utility

# Configure the NAT Utility

Configure the NAT utility on the Master server. Provide the IP address of the Control Service and Replica to connect all hosts and the Control Service.

**Follow these steps:**

1. On the Master server, navigate to the \CA\Arcserve RHA\Engine folder.

2. Locate and run the natutlgui.exe file to open the NAT Settings window.

3. Type the IP address and port number of the server where you have installed RHA Control Service.

4. Click Add to type the IP address and port number of the Replica Engine.

5. Click OK to let the NAT utility verify the connection to the Control Service and Replica. The status Connectable confirms the connectivity.

**Note:** To use the utility from the command line, use natutl.exe.

**Note:** If the Replica is using a NAT device, use the same process by configuring the NAT utility on the Replica rather than the Master.

## Create a Scenario using the NAT Utility

After you configure and test the connection from the NAT utility to the Replica and RHA Control Service, create a scenario as described in *Chapter 3: Creating Replication and High Availability Scenarios.*

# Chapter 4: Creating Replication and High Availability Cloud Scenarios

This section contains the following topics:

# Overview

Effective with this release, you can replicate and provide application high availability in the cloud. The *Full System Data Replication/High Availability to EC2* feature is an extension to the existing full system scenario type where Arcserve RHA enables replication and high availability of an entire Windows system into a VM running on Microsoft Hyper-V, Citrix Xen, or VMware ESX Hypervisor. The *Full System Data Replication/High Availability to EC2* feature extends the list of supported virtual environments by adding support for Amazon EC2.

To create a full system HA or replication scenario with an EC2 replica, ensure the following prerequisites are met:

- You must have an account with Amazon Web Services (AWS).

- You must have a configured VPC (Amazon Virtual Private Cloud) and have created the VPN connection between the on-premises network (where the master servers reside) and the VPC.

   **Note:** More information about Amazon VPC is described on the Amazon [website](website).
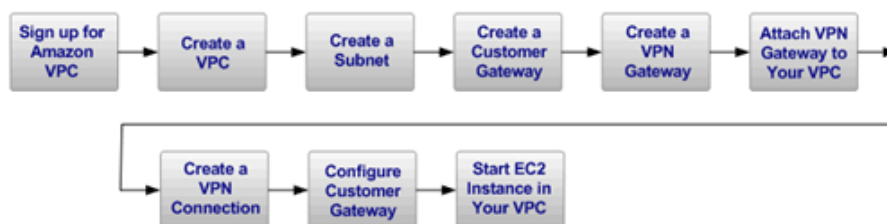


After these requirements are met, you replicate and provide application high availability by performing these steps and using these features:

- Use the Arcserve RHA Replication to Cloud option in the Scenario Creation Wizard.

   – Register the AWS account and credentials in Arcserve RHA

   – Retrieve the VPC ID associated to the account

   – Decide the EC2 instance and launch it

   – Remotely deploy the Arcserve RHA Engine

   – Use Arcserve RHA as usual

- New Arcserve RHA interface tab for cloud management (Cloud View).

   – Shows list of managed AWS accounts, instances, snapshots, EBS volumes, elastic IPs, security groups and so on.

# Full System EC2 Failover

For a Full System EC2 failover, the replica initiates the failover procedure if the master server becomes unresponsive. During failover the following actions occur:

- A new EC2 recovery instance of the same major OS version and processor architecture as the master is started using one of the predefined and supported AMIs. For example, if the Full System EC2 scenario protects a Windows 2008 x86_64 master server, then the Arcserve RHA appliance runs the new EC2 instance using Windows 2008 x86_64 AMI (ami-90d420f9). The following public AMIs will be used with the product:

    – ami-9ed420f7—Windows 2003 x86_64 recovery instance (us-east-1 region)

    – ami-98d420f1—Windows 2003 x86 recovery instance (us-east-1 region)

    – ami-ea45b883—Windows 2008 R2 x86_64 recovery instance (us-east-1 region)

    – ami-8bba8dff—Windows 2008 R2 x86_64 recovery instance (eu-west-1 region)

    – ami-61b98e15—Windows 2003 x86_64 recovery instance (eu-west-1 region)

    – ami-57b98e23—Windows 2003 x86 recovery instance (eu-west-1 region)

- The boot EBS volume of the recovery instance is detached from recovery instance and is attached to Arcserve RHA appliance.

- The necessary drivers are copied from the recovery instance's boot volume to the replicated master's boot volume, which is attached to the Arcserve RHA appliance. Additionally the replicated master's volume is mounted and the necessary registry entries are created to enable the drivers that were copied inside the replicated Windows image.

- The boot volume of the original recovery instance is detached from Arcserve RHAappliance and is deleted.

- Replicated EBS volumes are detached from Arcserve RHA appliance and are attached to the Full System EC2 recovery instance.

- The Full System EC2 recovery instance is started.

- All of the other standard failover steps are executed as usual. For example, if it is enabled, DNS redirection is executed.

The following illustration represents a Full System EC2 scenario protecting an on-premises server with two EBS volumes before failover occurs:
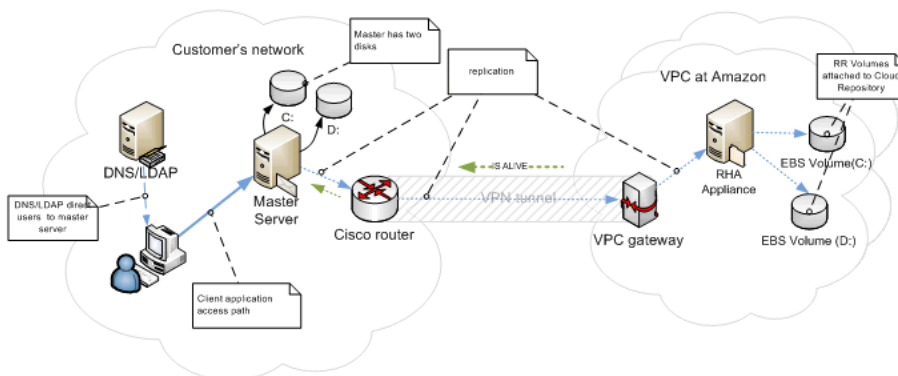


This illustration shows what happens once failover occurs:
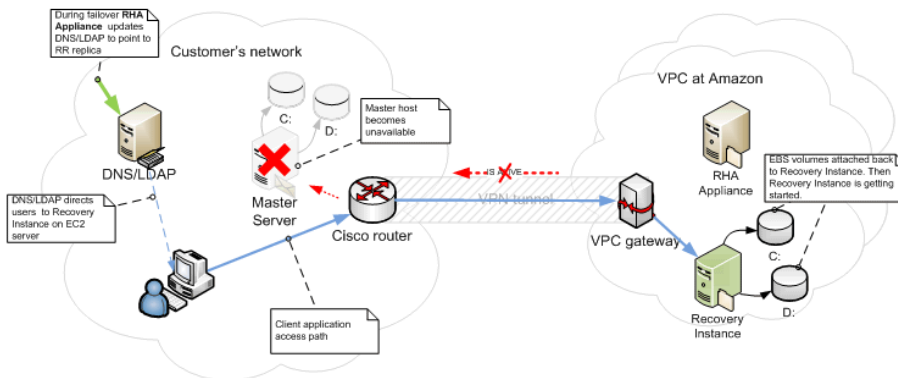
# Working with Arcserve RHA High Availability Cloud Scenarios

The *Full System High Availability to EC2* feature is an extension to the existing Full System scenario type where Arcserve RHA enables high availability of an entire Windows system into a VM running on Microsoft Hyper-V, Citrix Xen, or VMware ESX Hypervisor. The *Full System High Availability to EC2* feature extends the list of supported virtual environments by adding support for Amazon EC2.

Using the Scenario Creation Wizard to create a Full System High Availability EC2 scenario is the same as creating a non-cloud High Availability scenario, with the following exceptions:

- In the Master/Replica Hosts Assignments pane, after you have assigned the master host, select the *Amazon EC2* server type in the *Replica Settings* section. Select the EC2 replica instance (appliance).

- In the Master and Replica Properties screen, a new *Cloud* section is available. All of the properties are read-only except the *Clean up Cloud Resources When Remove Scenario* property that is disabled by default.

- In the Switchover Properties screen, from a connectivity perspective, VPC is a separate network and therefore the only redirection method that can be used for switchover is *Redirect DNS*. Currently all of the redirection methods are disabled by default, which means that during failover the EC2 replica appliance will create an EC2 recovery instance without redirecting traffic to it. If traffic redirection is required, then you have to enable the *Redirect DNS* option.

  **Notes:**

  - The *Physical Network Mappings* option designates the VPC subnetwork where the recovery instance will be started should Failover or Switchover occur. The Scenario Creation Wizard creates default mapping, but you can change the VPC subnets if needed.

  - Information about available VPC subnets is retrieved by CS from the EC2 replica appliance at scenario creation time. Then the EC2 replica appliance communicates with the EC2 API server, which can change depending on the VPC region in which it is located. This imposes the following requirements:

    - Appliance instances running in VPC should have access to the internet (your network administrator should enable the corresponding routing rules).

- ◆ Appliance instances should have a proper DNS server set so the IP addresses of EC2 API servers can be resolved.

- For more information about creating high availability scenarios, see Creating Replication and High Availability Scenarios.

# Working with Arcserve RHAReplication Cloud Scenarios

The *Data Replication to EC2* feature is an extension to the existing scenario types where Arcserve RHA enables replication of an entire Windows system into a VM running on Microsoft Hyper-V, Citrix Xen, or VMware ESX Hypervisor. This feature extends the list of supported virtual environments by adding support for Amazon EC2.

Using the Scenario Creation Wizard to create a Data Replication EC2 scenario is the same as creating a non-cloud Data Replication scenario, with the following exceptions:

- In the Select Server and Product Type screen, Microsoft Hyper-V is currently not supported for cloud-based replication scenarios.

- In the Master/Replica Hosts screen, after you have assigned the master host, use the Replicate to Cloud check box and click the Select Cloud Host button so you can specify Amazon EC2 as the server, and select an EC2 replica instance (appliance).

- In the Master and Replica Properties screen, a new *Cloud* section is available. All of the properties are read-only except the *Shutdown Instance on scenario stop* property that is disabled by default.

**Note:** For more information about creating replication scenarios, see Creating Replication and High Availability Scenarios.

# Replicate to Cloud

You can now protect an on-premises server by replicating it to an Amazon Web Services (AWS) cloud-based target. The Arcserve RHA Scenario Creation Wizard lets you create and configure a VPC and a VPN connection between your on-premises network and a cloud-based target.
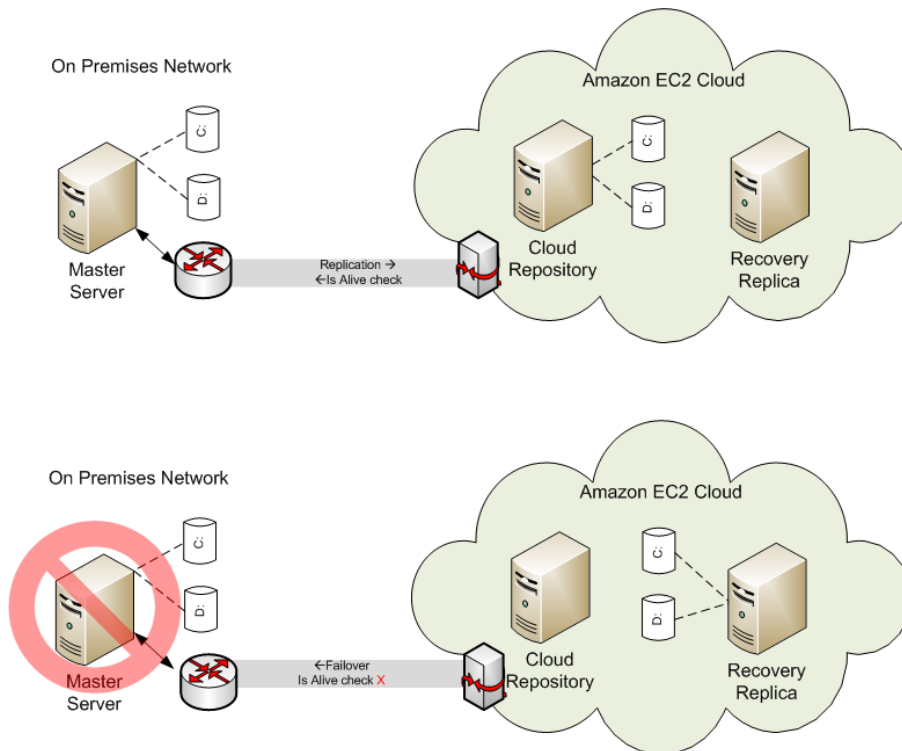
- The Cloud Repository server is an EC2 instance with Arcserve RHA installed.

- The Recovery Replica is also an EC2 instance with the same disk layout as the Master server. After the Recovery Replica is created, it is stopped. All volumes are detached from it and attached to the Cloud Repository instance. Arcserve RHA scenarios are created from on-premises Master servers to the exposed volumes on the Cloud Repository.

- If Is-Alive fails and failover is triggered, the exposed volumes attached to the Cloud Repository are attached back to the Recovery Replica. The Recovery Replica is then started. Users are routed to this server until switchback occurs.

# Navigating the Cloud View Panel

The Cloud View panel lets you manage and monitor your EC2 instances and supports multiple AWS accounts.

**Note:** If you need to manage EBS volumes, snapshots, Elastic IPs, security groups, and so on, use the standard AWS management tools such as the AWS Management Console.

This panel provides the following options:

**Cloud View**

Displays the cloud accounts and instances. The Cloud View tab contains two panes. The left pane displays the cloud accounts and instances, including instance, status, key pair name, security group and IP address.



**AWS Cloud Account**

Displays the AWS statistics panel that provides the EC2 resource use summary for a cloud account. This pane displays on the right when an AWS account entry is selected in the left pane.
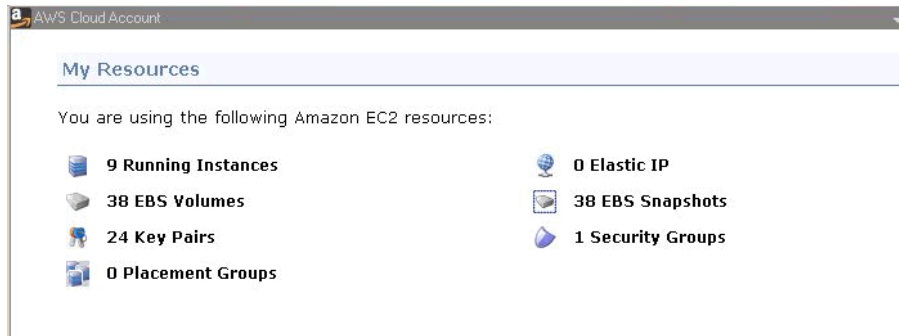
## AWS Cloud Instance

Displays the AWS statistics panel with detailed information about an instance. This pane displays on the right when a specific instance is selected in the left pane.
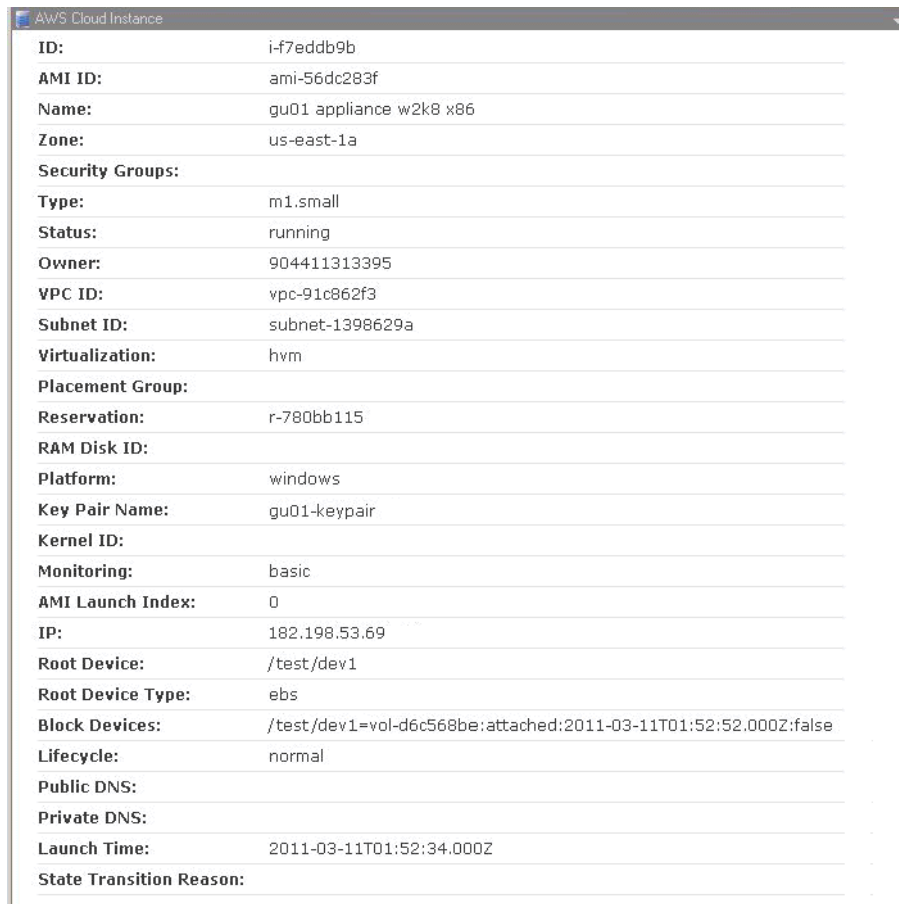


## AWS Cloud Watch

Displays additional instance statistics when the AWS Cloud Watch option is enabled.
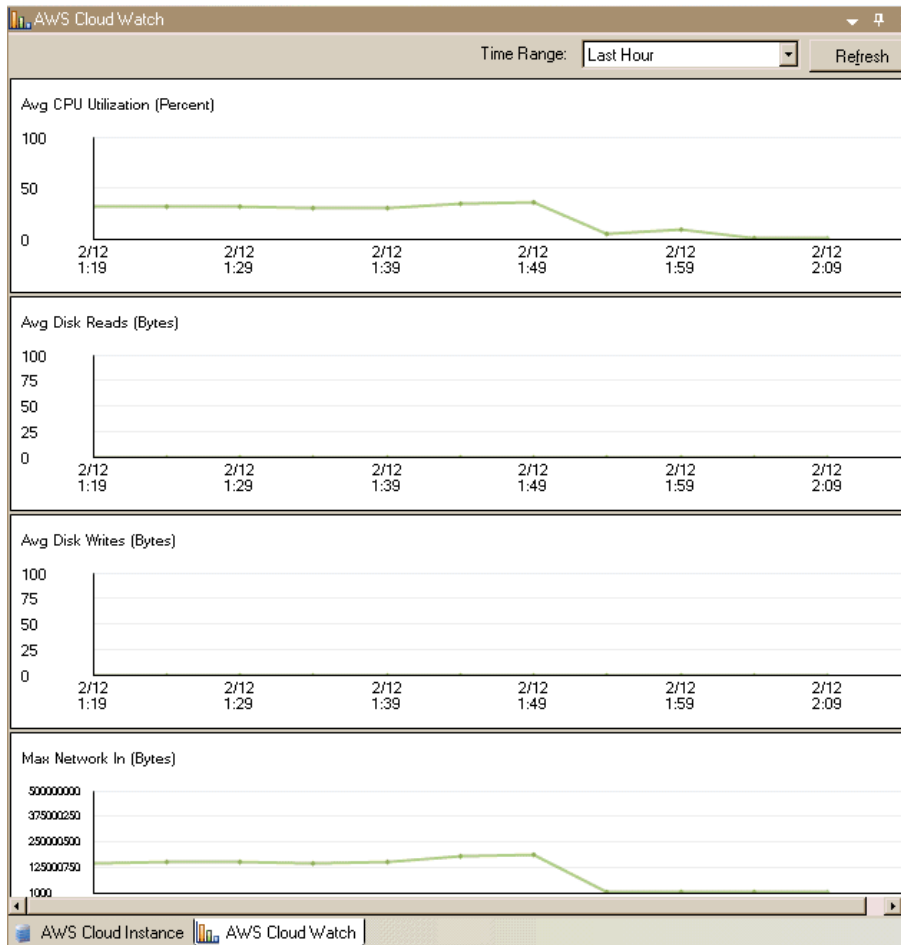
Using the Cloud View panel activates the Cloud View Toolbar. Use the following options on the Cloud View toolbar to help you easily manage your cloud accounts and instances:

**Add Cloud Account**

   Opens the Add Cloud Account dialog so you can create a new cloud account.

**Delete Cloud Account**

   Specifies to delete the selected cloud account.

**Update Cloud Account**

   Lets you change cloud account information.

**Change Default Region**

   Lets you change the AWS default region.

**Create New Instance**

   Opens the Instance Creation Wizard so you can create a new instance.

**Delete Instance**

Specifies to delete the selected instance.

**Start Instance**

Specifies to start the selected instance.

**Stop Instance**

Specifies to stop the selected instance.

**Reboot Instance**

Specifies to restart a stopped instance.

**Get Password**

Lets you retrieve an instance password.

**Refresh**

Refreshes the cloud account and instances that display in the Cloud View.
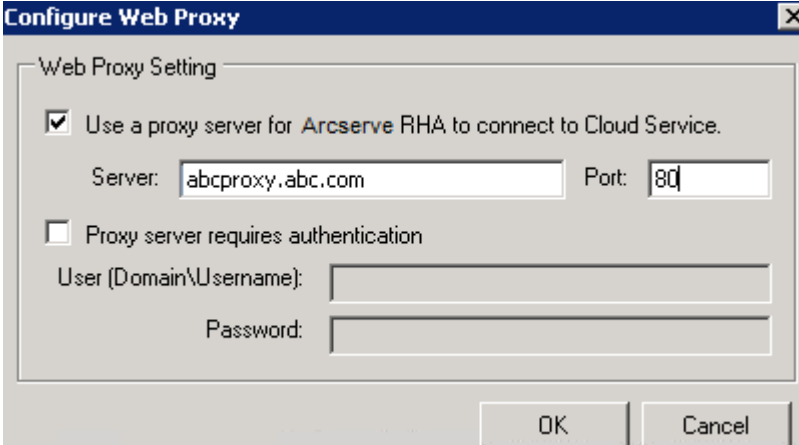
**Configure Web Proxy Server**

Lets you specify proxy information such as IP address, port, and user credentials so that you can use a proxy to connect to the cloud service.

# Configure the Web Proxy to Connect to the Cloud Service

If you want to use the *Use proxy to connect to the cloud service* option in the Add Cloud Account dialog, you must first configure the Web proxy that you want to use to manage EC2 resources.

To configure the Web proxy, click the Configure Web Proxy 🌐 toolbar button in the Cloud View panel and enter web proxy setting information (such as server, port, and user credentials). A test request is sent to the server to verify the proxy setting. Once verified, the proxy setting is saved in the AWS account.

# Add a New Cloud Account

To see and manage EC2 instances in the Cloud View panel, you must first add a new cloud account using your AWS credentials.

**To add a new cloud account**

1.  Click the Add Cloud Account button 🖥️ on the toolbar.

    The Add Cloud Account dialog opens.
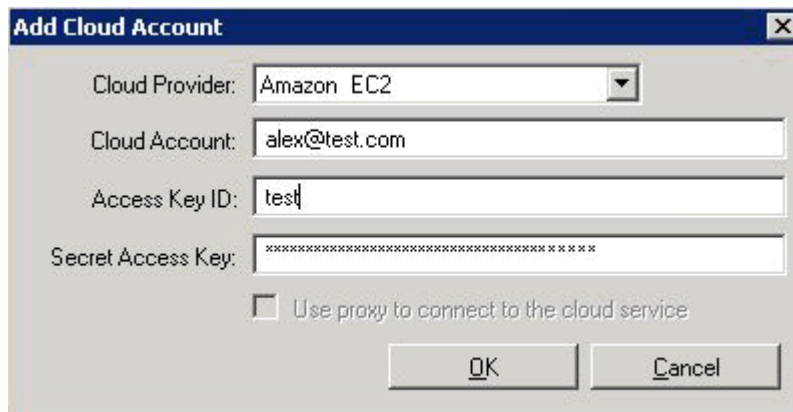


2.  Enter information in the following fields:

    **Cloud Provider**

    Specifies the name of the cloud provider service.

    **Cloud Account**

    Specifies the cloud account; typically this is the email address you used to register your AWS account.

    **Access Key ID**

    Lets you enter the access key ID for your AWS account.

    **Secret Access Key**

    Lets you enter the secret access key provided by your AWS account.

    **(Optional) Use proxy to connect to the cloud service**

    Lets you specify whether to use a Web proxy to communicate with the AWS web services server. If you select this check box to enable this option, ensure that you have first configured this proxy.

    **Note:** The information required for these fields can be obtained from your AWS account home page using the Security Credentials tab, as shown here:

## Access Credentials

There are three types of access credentials used to authenticate your requests to AWS services: (a) access keys, (b) X.509 certificates, and (c) key pairs. Each access credential type is explained below.

| Access Keys | X.509 Certificates | Key Pairs |

Use access keys to make secure REST or Query protocol requests to any AWS service API. We create one for you when your account is created — see your access key below.

**Your Access Keys**

| Created | Access Key ID | Secret Access Key | Status |
|---------|---------------|-------------------|--------|
| January 27, 2009 | OO3KNR20D32SJNAQ5ET2 | Show | Active (Make Inactive) |
| October 11, 2010 | A5ETY8A6DJS2A92NSKA6 | Show | Active (Make Inactive) |

View Your Deleted Access Keys

For your protection, you should never share your secret access keys with anyone. In addition, industry best practice recommends frequent key rotation.

Learn more about Access Keys

3. Click OK.

The Add Cloud Account dialog closes. The cloud account displays in the Cloud View panel as a registered cloud account, and it displays all of the EC2 instances that belong to that account.

# Change the Default AWS Region

In the Cloud View panel, you can select a different AWS region while you are managing EC2 instances. To change the default AWS region, click the Change Default Region button  on the toolbar. In the Change Default Region dialog, select a different region from the drop-down list. The Cloud View panel display is updated with the available instances for the region you selected.

# Update Cloud Account Information

You can update the credentials for a previously configured cloud account. For example, if the Access Key ID and Secret Access Key were changed (a new pair was generated and the previous pair was deactivated) using the Amazon Management Console, the AWS account credentials must be manually updated. To update cloud account credential information, select the cloud account in the Cloud View panel that you want to update and click the Update Cloud Account button  on the tool-bar. Enter the new credentials and click OK. The cloud account information is updated in the Cloud View pane.

# Delete a Cloud Account

You can delete a cloud account that you no longer use. To delete a cloud account, select the cloud account in the Cloud View panel that you want to delete and click the Delete Cloud Account button ![icon] on the toolbar. The cloud account is removed from the list in the Cloud View panel.

# Create a New EC2 Replica Instance

To use Arcserve RHA cloud capabilities for a Full System Data Replication or High Availability scenario, you need to create an EC2 replica instance. Ensure the following requirements are met before starting this procedure:

- Amazon Web Services (AWS) account is created

- The Amazon Virtual Private Cloud (VPC) is created and connected with the on-premises network using VPN

  **Note:** More information about Amazon VPC is described on the Amazon website.

- The AWS account is registered in the Arcserve RHA Manager

**Note:** More information about EC2, including instance details and creating key pairs (required for this procedure), is described in the Amazon EC2 user documentation on the Amazon website.

**To create a new EC2 replica instance**

1. Select the Cloud View panel in the Arcserve RHA Manager.

   Instances that you have already created display in the Cloud Accounts and Instances list for the specified region. You can click the Change Default Region button 🌐 on the toolbar to specify a different default region.

2. Right-click the AWS account that you want to use, and select Create Instance.

   The Instance Creation Wizard opens.
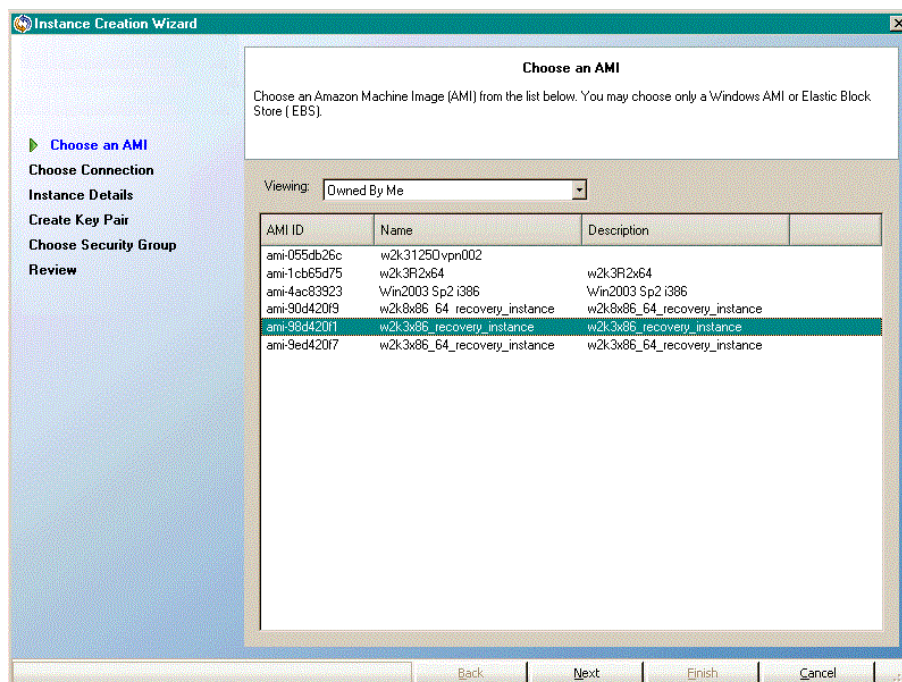
3. Select an Amazon Machine Image (AMI) from the list in the Choose an AMI dialog and click Next.
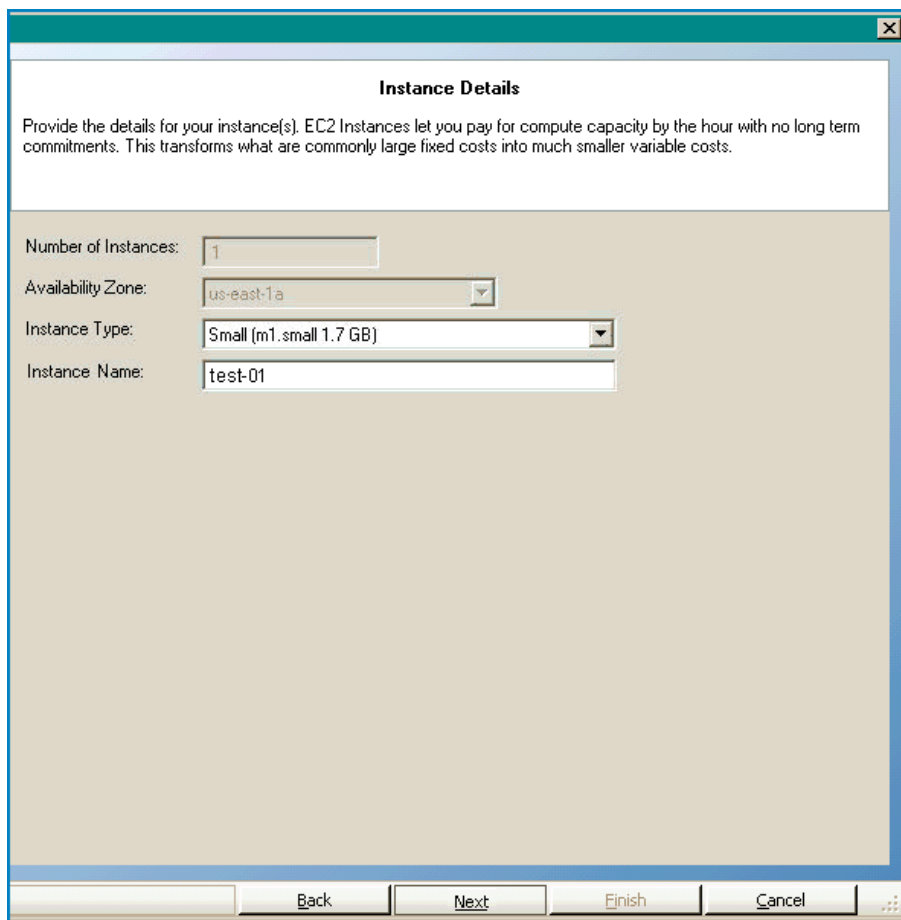
   **Note:** You can only use EBS-backed Windows AMIs.

   The Choose Connection screen opens.

4. Specify the VPC subnet where you want to allocate the instance and click Next.

   **Note:** The instance is automatically assigned a DHCP address from the IP range of the subnetwork where it was allocated.

   The Instance Details screen opens.



5. Specify the number of instances, Availability Zone, and Instance Type and click Next. The instances let you pay for compute capacity by the hour with no long-term commitments so that large fixed costs are transformed into much smaller variable costs.

   The Create Key Pair screen opens.
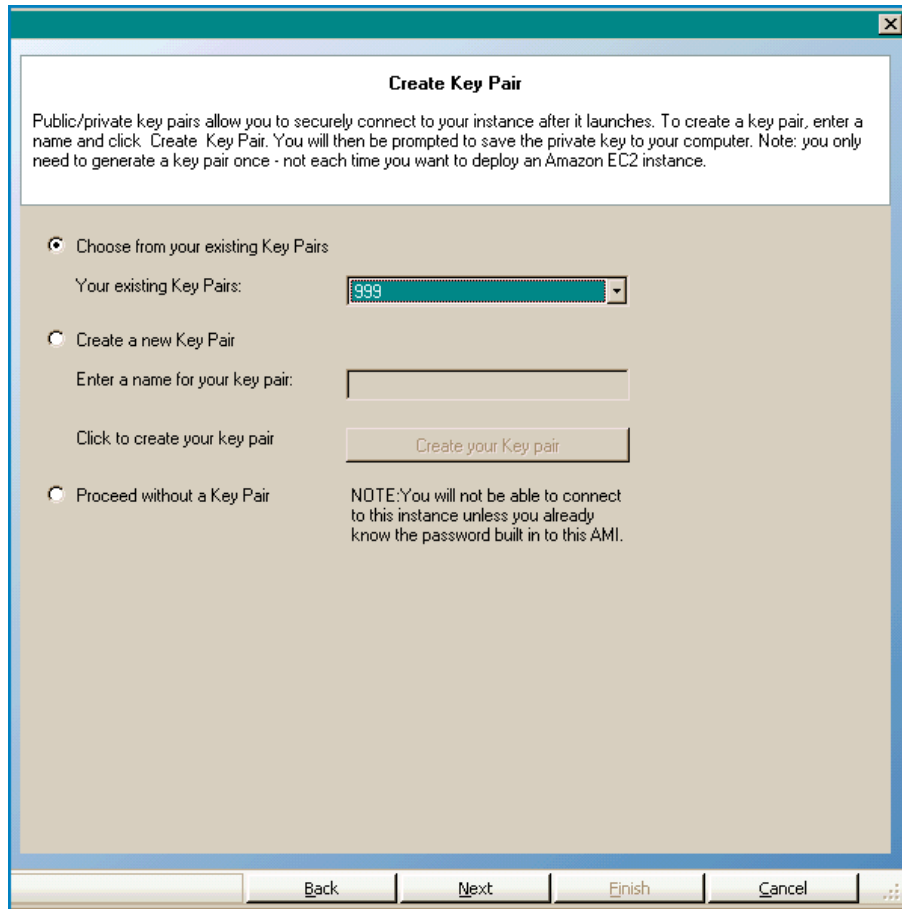
6. Select an existing key pair or create a new key pair to securely connect to your instance after it launches. If you create a new key pair, you are prompted to save the private key to your computer. Click Next.

   **Note:** You only need to generate a key pair once, not each time you want to deploy an instance.

   The Choose Security Group screen opens.
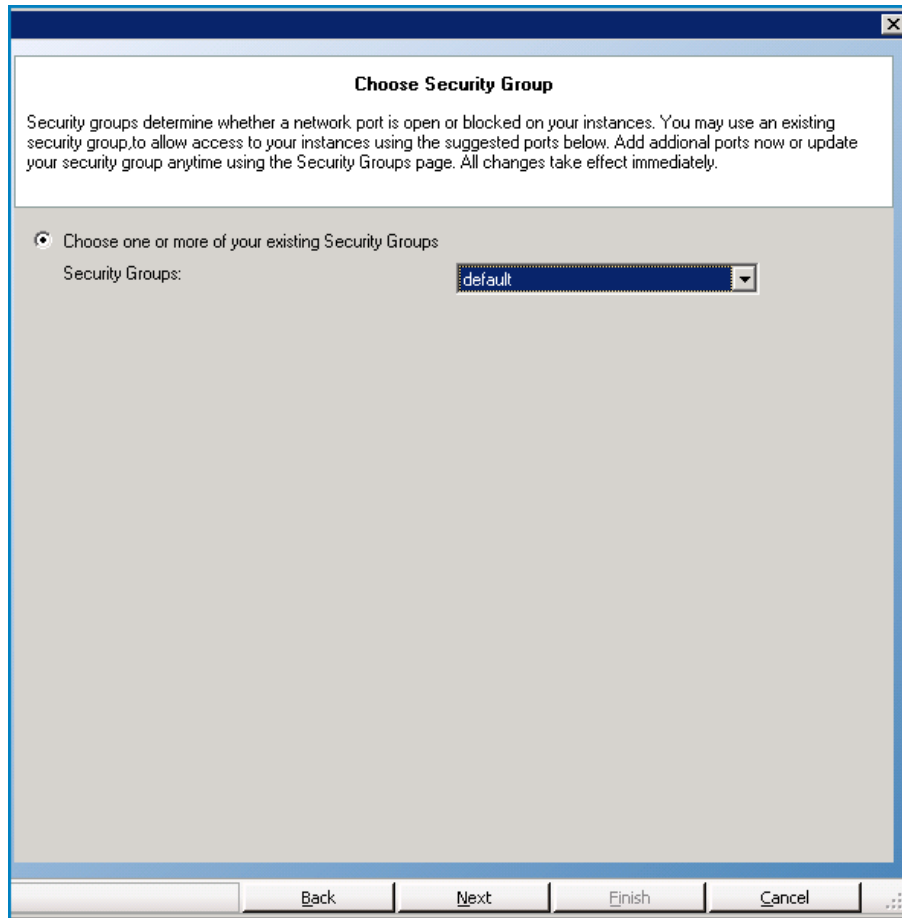
7. Select a security group and click Next.

   The list shows a list of existing security groups.

   The Review screen opens.

**Review**

Please review the information below, then click Finish.

## Review

| | |
|---|---|
| **AMI ID** | ami-98d420f1 |
| **AMI Name** | w2k3x86_recovery_instance |
| **Description** | w2k3x86_recovery_instance |
| **Instance Name:** | test-01 |
| **Number of Instances:** | 1 |
| **Availability Zone:** | us-east-1a |
| **Instance Type:** | Small (m1.small 1.7 GB) |
| **Key Pair Name** | 999 |
| **VPC Subnet** | 192.168.24.0/51 |

Back    Next    Finish    Cancel

8. Review the information you have specified and click Finish.

The instance is created and you will receive a default password.

**Notes:**

◆ The Cloud View panel only allows creation of VPC-based instances. You cannot allocate an instance outside VPC, so VPC configuration is mandatory.

◆ You can create non-VPC instances (public instances) using the AWS management tools (in the AWS Management Console). These public instances will be displayed in the Cloud View panel. However, those instances will not be available in the Scenario Creation Wizard because the wizard only supports "private to private" scenarios, or replication from a private on-premises network to VPC.

◆ You can filter resource browsing by selecting different regions. There are seven regions available to AWS users: EU West (Ireland), US East (Virginia), US West (N. California), US west (Oregon), South America (Sao Paulo), Asia Pacific (Tokyo), and Asia Pacific (Singapore). Currently AWS only allows creation of one VPC per AWS account. Every VPC may have

several subnetworks associated with it (you select the subnetwork ID when the instance is allocated). The subnetwork may reside in one of the four regions. If you want to allocate an instance in a specific subnetwork, then you should first select the corresponding region in the *Change Default Region* drop-down, however, note that RHA only supports five regions: US East (Virginia), EU West (Ireland), Asia Pacific (Tokyo), Asia Pacific (Sydney), and South America (Sao Paulo).

# Start an EC2 Replica Instance

To start an EC2 instance in the Cloud View panel that was stopped, select the EC2 instance that you want to start and click the Start Instance button  on the toolbar. The status of the selected EC2 instance in the Cloud View panel changes from *Stopped* to *Run Pending*, and finally to *Running*.

# Stop an EC2 Replica Instance

To stop (shut down but not remove) an EC2 instance in the Cloud View panel that is no longer used, select the EC2 instance that you want to stop and click the Stop Instance button  on the toolbar. The status of the selected EC2 instance in the Cloud View panel changes from *Running* to *Stopped*.

# Delete an EC2 Replica Instance

To delete an EC2 instance in the Cloud View panel that is no longer used, select the EC2 instance that you want to delete and click the Delete Instance button  on the toolbar. The deleted EC2 instance is removed from the instance list in the Cloud View panel.

# Create a New Full System EC2 High Availability Scenario

You can create a Full System EC2 High Availability scenario and replicate an entire on-premises Windows system into offline EBS-backed Windows AMIs that will be brought online should the master server fail. This procedure launches a Wizard that guides you through the steps required for High Availability scenario creation. However, properties can also be configured outside of the wizard.

**To create a new full system EC2 High Availability Scenario**

1.  Open the Manager and choose Scenario, New or click the New Scenario button to launch the wizard.

    The Welcome screen opens.

2.  Choose Create a New Scenario, select a Group from the list and then click Next.

    The Select Server and Product Type screen opens.
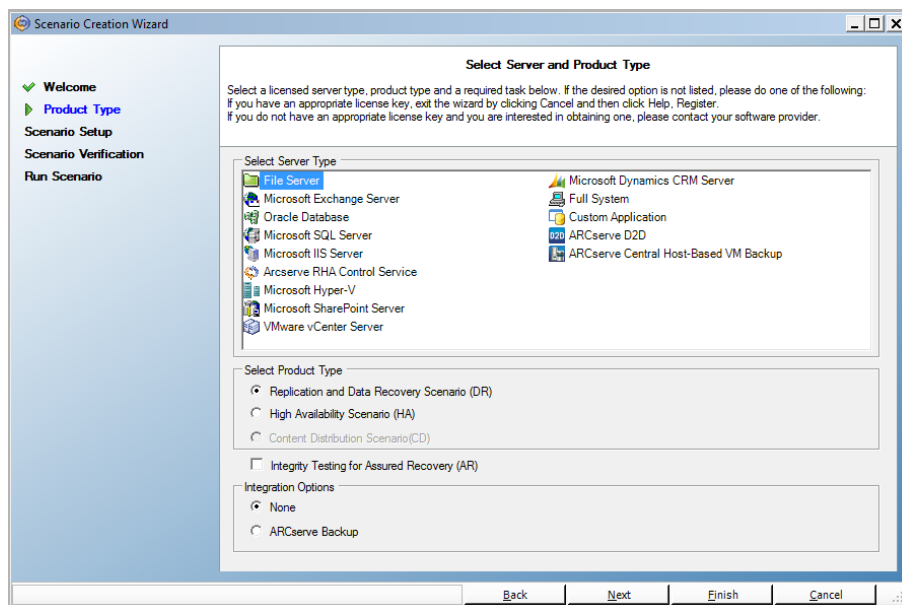


3.  Choose Full System, High Availability Scenario (HA) and then click Next.

    The Master and Replica Hosts screen opens.
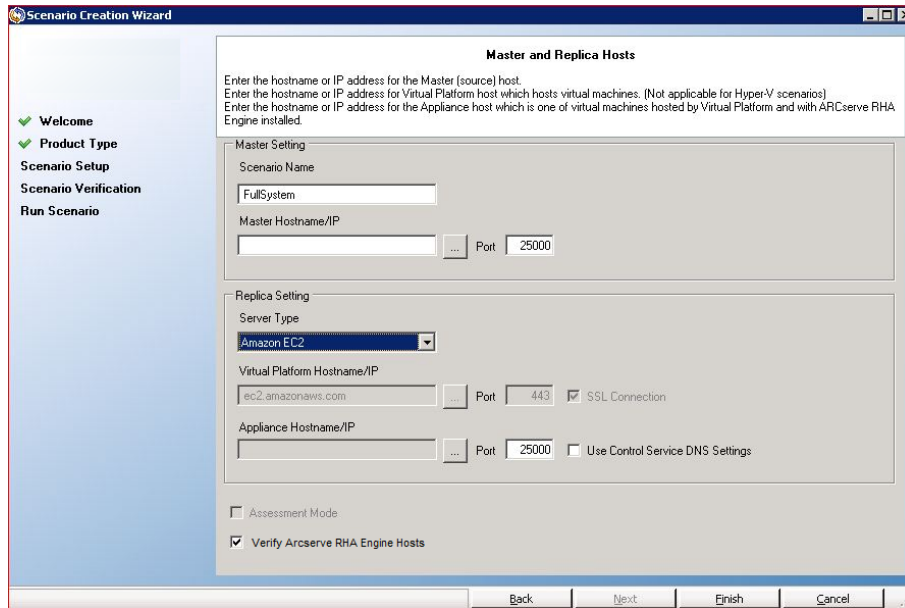
4. Do the following in the Master and Replica Hosts screen:

   a. Type a Scenario Name and enter the Hostname or IP Address and Port number for the Master server.

   b. Specify Amazon EC2 as the Replica server.

   c. Specify the EC2 replica instance (appliance). Click the [ ··· ] button to browse for and select the AWS account and EC2 replica instance (appliance).

   The Cloud Instance Selection dialog opens.

   d. Select the AWS account, cloud replica (appliance), and region and click OK.

   e. Select or clear the Use Control Service DNS check box as needed. Selecting this check box applies the DNS settings from the Control Service server to the EC2 replica instance (appliance) host.

   f. Ensure the Verify Arcserve RHA Engine on Hosts option is enabled (default) and click Next.

The Engine Verification screen opens.
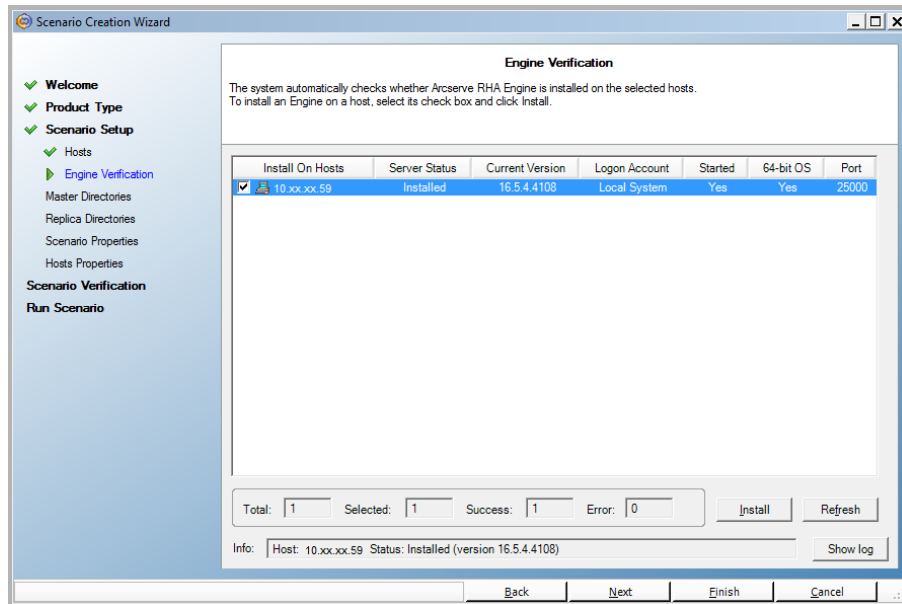
5. Wait for Engine Verification to complete and click Next. If required, click Install to upgrade the Engine on one or both servers and Verify Again.

   The Volume Setting screen opens.

6. Select one or more volumes for the physical machine you want to protect and click Next.

   The Scenario Properties screen opens.

7. Accept the default values or set new values as required and click Next when you are done.

   **Note:** Scenario properties control the entire scenario. These properties can also be configured outside of the Wizard. For more information, see Configuring Scenario Properties.

   **Note:** The Network Adapter Mapping dialog opens if there are more than one replica network adapters.

   The Master and Replica Properties screen opens.
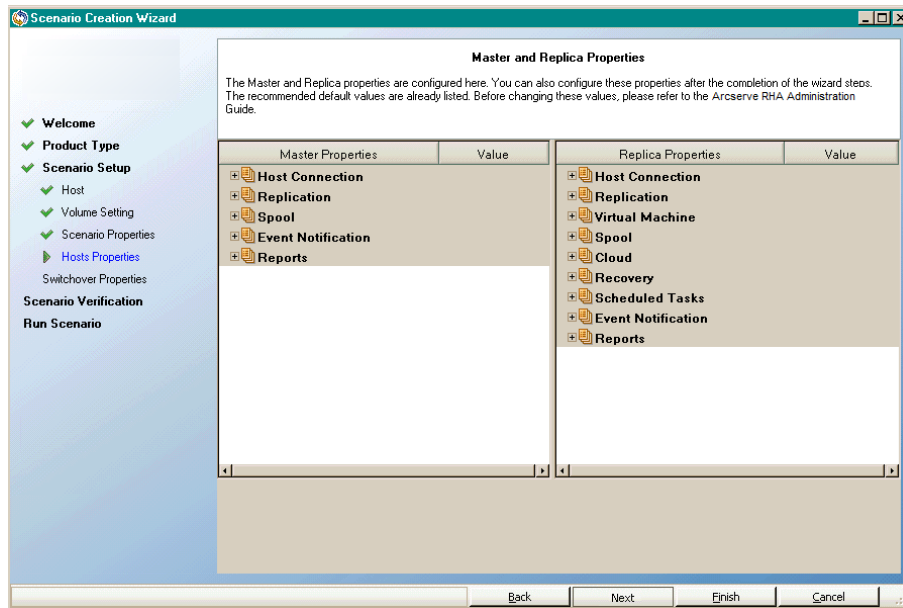
8.  Master and Replica properties apply only to host servers. Accept the default values or change values as needed and click Next.

**Note:**

◆ All Cloud Replica properties are read-only except for the *Cleanup Cloud Resources When Remove Scenario* property, disabled by default.

◆ General purpose (SSD) volume, provisioned IOPS (SSD) volume, and mag-netic volume can be selected as the EBS volume type.

◆ EBS volume encryption can be enabled for the specified EC2 instance types. For more information on the supported EC2 instances types, refer the following link: http://-docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html.

◆ Network mapping must be assigned for Assure Recovery, otherwise the assure recovery will fail.

To modify the network mapping, expand the Physical Network Mappings option under Virtual machine.
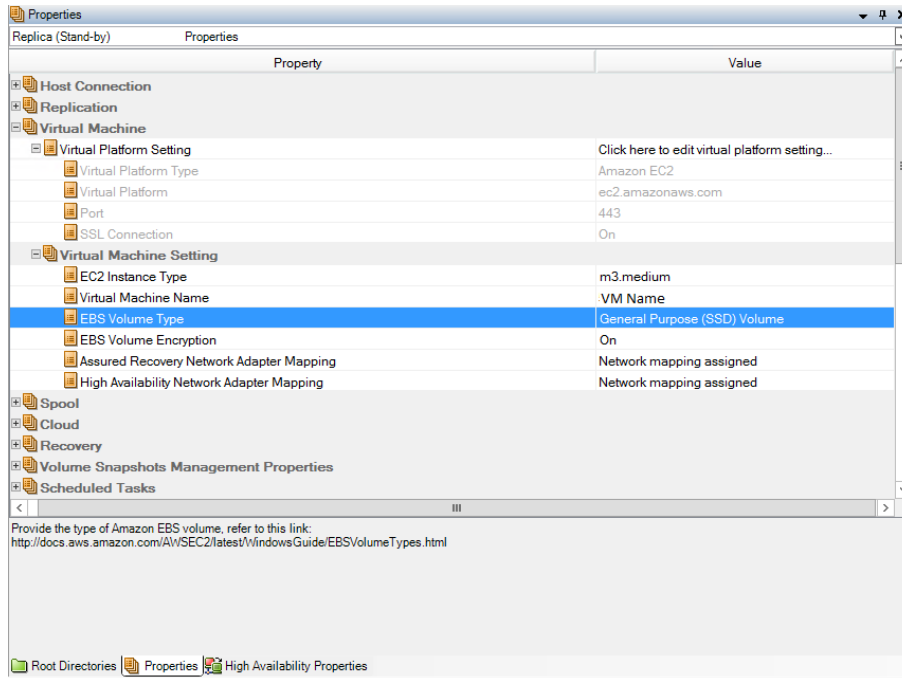
Click "Assured Recovery Network Adapter Mapping" or "High Availability Net-work Adapter Mapping".

The Network Adapter Mapping dialog opens.
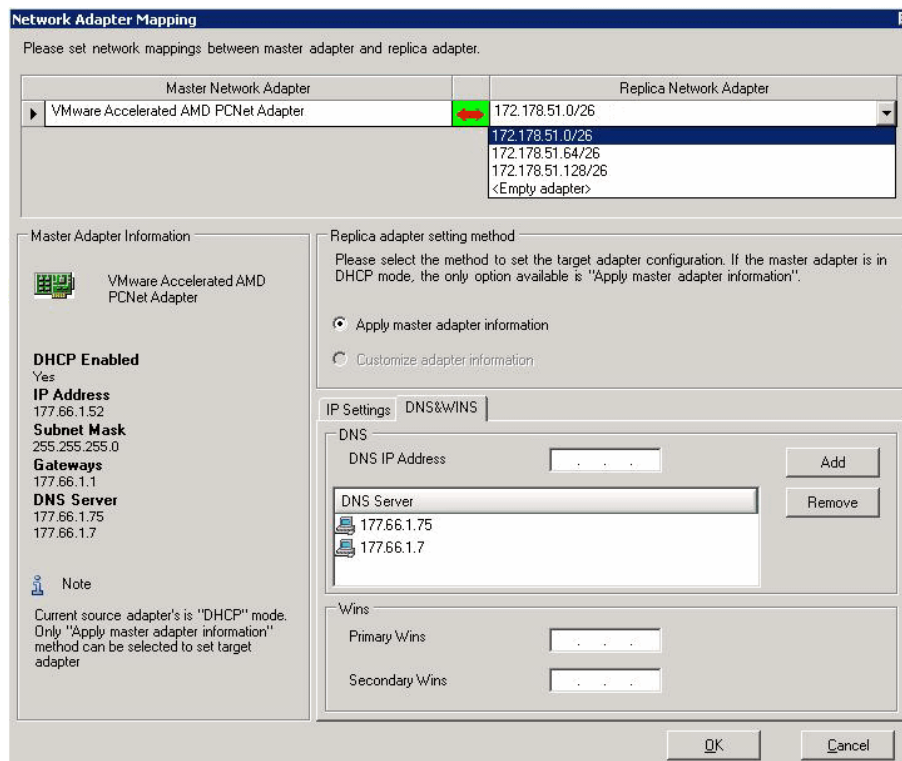


Set network mapping between the master adapter and the replica adapter and click OK.

The Master and Replica Properties opens.

Click Next.

The Switchover Properties screen opens.

9. Expand the *Switchover* property and enter the Switchover Hostname.

10. Expand the *Hosts* property and enter the Master Fully Qualified Name and Replica Fully Qualified Name.

11. Expand the *Network Traffic Redirection* property and specify redirection options, including Redirect DNS, DNS Servers IPs, and Master IPs in DNS.

   **Note:** When you set the Redirect DNS option to *Off*, you can also specify a value for the Virtual Machine IPs on Replica server in DNS option. If the Redirect DNS property value is *On*, then the Virtual Machine IPs on Replica server in DNS option will not display in the list.

   The Switchover and Reverse Replication Initiation screen opens.
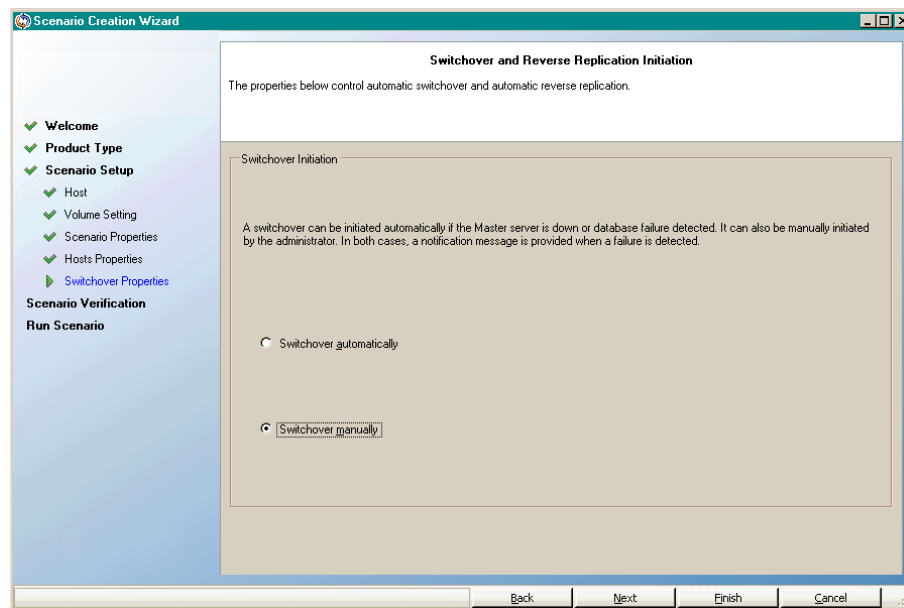


12. Select a switchover option. Automatic is not recommended. For more information, see Switchover. Click Next.

   Wait for Scenario Verification to complete.

13. If errors or warnings are listed, resolve them before continuing. When ready, click Next.

   The Scenario Run screen opens.

14. Click Run Now to start synchronization and activate the scenario, or click Finish to run the scenario later.

# Create a New EC2 Data Replication Scenario

You can create an EC2 Data Replication scenario allowing EC2 instances specified in the Scenario Creation Wizard to be used as Replica servers. This procedure launches a Wizard that guides you through the steps required for Data Replication scenario creation. However, properties can also be configured outside of the wizard.

**To create a new EC2 Data Replication scenario**

1. Open the Manager and choose Scenario, New or click the New Scenario button to launch the wizard.

   The Welcome screen opens.

2. Choose Create a New Scenario, select a Group from the list and then click Next.

   The Select Server and Product Type screen opens.
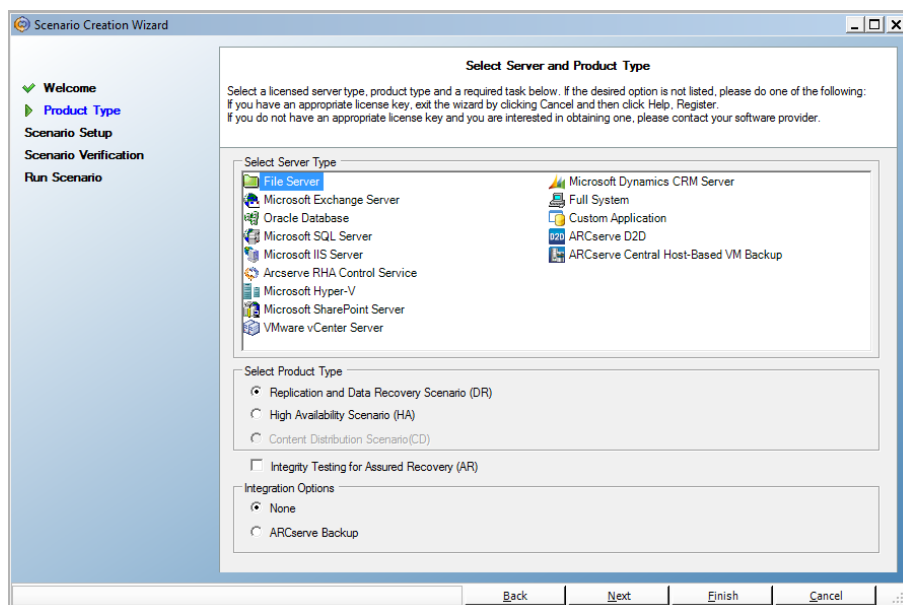


3. Select the Server Type and select Replication and Data Recovery Scenario (DR) and click Next.

   **Note:** Microsoft Hyper-V is not currently supported for cloud-based data replication.

   The Master and Replica Hosts screen opens.
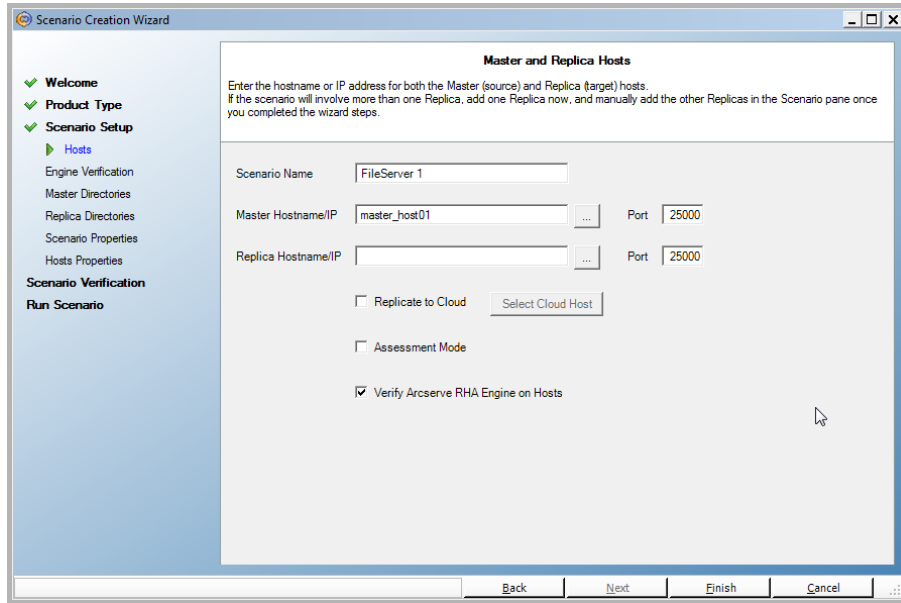
4.  Type a Scenario Name and enter the Hostname or IP Address and Port number for the Master server. Specify Amazon EC2 as the Replica server. Select the Replicate to Cloud check box and then click the Select Cloud Host button to specify the EC2 replica instance (appliance). Make sure the Verify Arcserve RHA Engine on Hosts option is enabled (default) and click Next.

The Cloud Instance Selection dialog opens.

5.  Enter the Cloud Provider, Cloud Account, and Region information and click OK.

**Note:** Click the Refresh button to refresh the list of instances.

The Engine Verification screen opens.
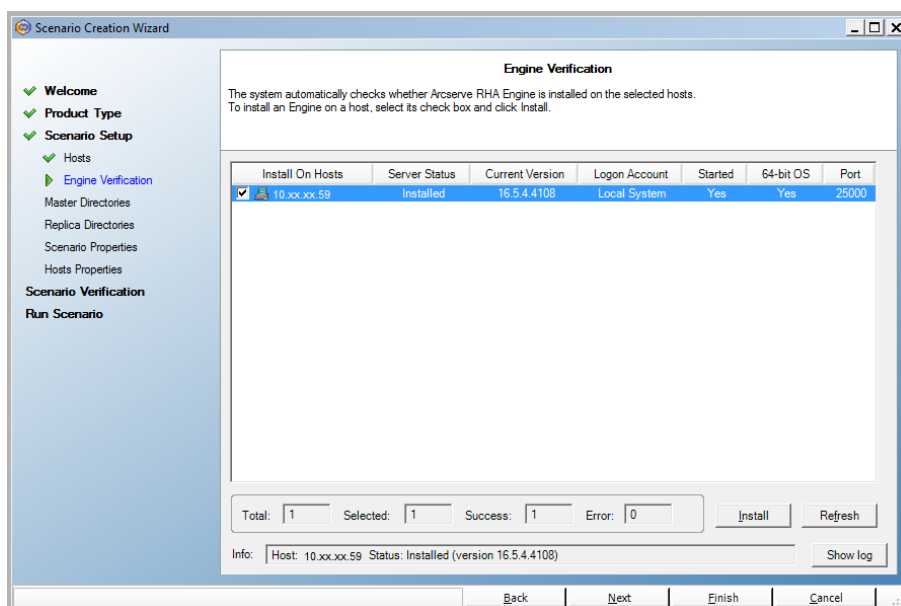
6. Wait for Engine Verification to complete and click Next. If required, click Install to upgrade the Engine on one or both servers and Verify Again.

The Volume Setting screen opens.

7. Enter information and click Next.

The Scenario Properties screen opens.

8. Accept the default values or set new values as required and click Next when you are done.

**Note:** Scenario properties control the entire scenario. These properties can also be configured outside of the Wizard. For more information, see Configuring Scenario Properties.

The Master and Replica Properties screen opens.
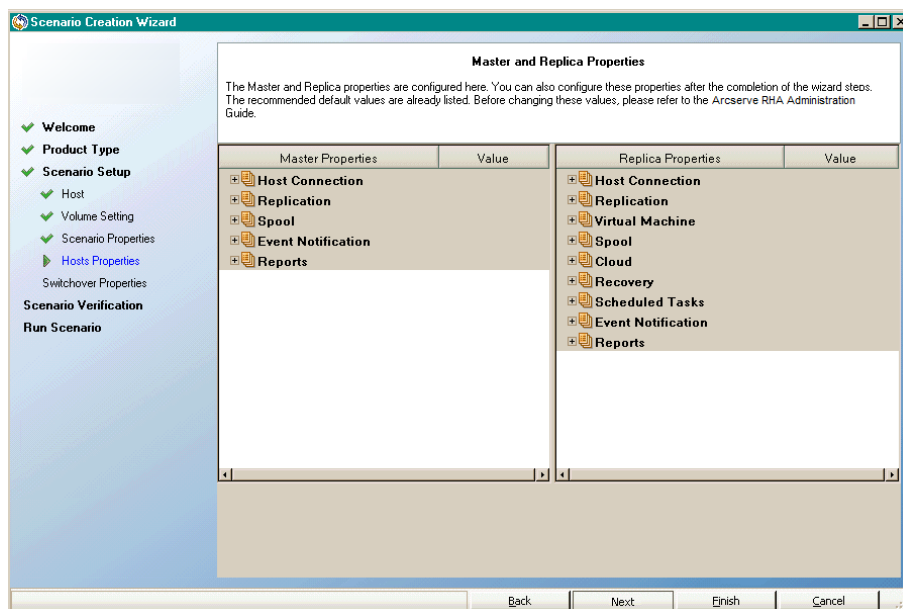


9. Master and Replica properties apply to only host servers. Accept the default values or change values as needed and click Next.

**Note:** All Cloud Replica properties are read-only except for the *Shutdown Instance on Scenario Stop* property, disabled by default. For more information see Cloud Properties.

10. Click Next.

Wait for Scenario Verification to complete.

11. If errors or warnings are listed, resolve them before continuing. When ready, click Next.

The Scenario Run screen opens.

12. Click Run Now to start synchronization and activate the scenario, or click Finish to run the scenario later.

# Run and Synchronize a Full System EC2 High Availability or Data Replication Scenario

After you create a scenario, you need to run it to start the replication process. A previously created Full System EC2 High Availability or Data Replication scenario is executed as any other Arcserve RHA scenario, with the following exceptions:

- For every replicated master volume the Arcserve RHA appliance creates and attaches an EBS volume of the same size.

- Every replicated master volume is initially synchronized with its corresponding EBS volume on the Arcserve RHA appliance (you can select volume, block or file synchronization).

- The replication flow is unified for all Full System scenarios (replication is at file system level, and file system changes are applied to the mounted EBS volumes).

**Note:** For information about the replication process, see Running the Replication Process.

# Performing Switchover for a Full System EC2 High Availability Scenario

For a Full System EC2 High Availability scenario, you can execute one of the following methods if the master server becomes unresponsive:

- Manual switchover
- Automatic failover

Switchover can be triggered automatically by Arcserve RHA when it detects that the Master is unavailable (failover). Alternatively, Arcserve RHA can simply alert you to the problem, and then you can manually initiate switchover from the Manager. Once triggered, whether manually or automatically, the switchover process itself is fully automated.

You select manual switchover, otherwise, if the automatic failover option is enabled and the master server is unresponsive, automatic failover occurs. The switchover process for a Full System EC2 High Availability scenario is identical to non-cloud scenarios, with the following exceptions:

- The EC2 RHA Appliance runs and puts the Failover EC2 instance into a Stop state using the AMIs prepared previously (Arcserve RHA provides four public AMIs). The Failover EC2 instance is instantiated from the AMI of the same major operating system version and processor architecture (W2K3 or W2K8 and x86 or X86_64).

- The boot volume of the Failover EC2 instance is detached and attached to the EC2 RHA appliance.

- Xen driver injection into the replicated volume is performed.
    - Xen drivers from the original boot volume of the Failover replica are copied.
    - The driver injection routine, which mounts and changes the registry on the replicated master volume, is executed.

- The original Failover EC2 Instance boot volume is detached and deleted from the machine from which the Xen drivers were copied.

- The master's replicated volumes are detached and then attached to the Failover EC2 Instance.

- The Failover EC2 Instance is started.

**Note:** For more information about switchover, review the topics in Switching Over and Switching Back.

# Recovery Using an EC2 Failover Replica

If you replicated your on-premises system to an EC2 replica instance and either manual switchover or automatic failover occurred, you can perform data recovery using the EC2 failover replica instance:

- Replicate the Full System EC2 instance to another virtual environment (such as the on-premises Xen/Hyper-V/ESX or to another EC2 RHA appliance)
- Restore individual data sets using the File System replication scenario

The recovery process using an EC2 failover replica is identical to non-cloud scenarios, with some differences. When switchover or failover occurs and a new EC2 failover instance is started, a backward replication scenario with the EC2 appliance as the new master and the on-premises host as the new replica is created:

- An automatic or manual backward scenario is run, which allows only block and file level synchronization.
- It excludes Windows system files from synchronization/replication.
- Lets you change replication directories/registry hives before the backward scenario is run.
- The registry is replicated.

**Note:** For information about the recovery process, see Recover Lost Data from a Replica.

# Chapter 5: Running the Replication Process

This section contains the following topics:

# Initiate Replication

After you create a scenario, you need to run it to start the replication process. Normally, before data changes on the Master can be replicated on the Replica, the Master and the Replica need to be synchronized. Therefore, the first step in initiating 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.

Arcserve RHA allows simultaneous synchronization and replication. In this case, the servers are synchronized while files are in use and being updated. All of these changes that occur on the Master are captured and held in a spool. After the synchronization is complete, the replication starts and the changes are updated on the Replica.

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

**To initiate replication**

1.  From the Manager, select the scenario you want to run.

2.  To run the scenario, click **Run** on the Standard toolbar.

    Arcserve RHA verifies the scenario before running it.

3.  If the scenario was not set up correctly or problems occurred in the participating hosts, errors are reported on the Event pane.

    Be aware of the following:

    - If any errors are displayed, you cannot run the scenario. These errors must be corrected before you can start the replication process.

    - Replication of mount points will succeed 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.
      This rule also applies to the removal of mount points during replication. If you want to remove mount points that you previously defined as part of the replication, do the following: first stop the Engine, then remove the mount points, and finally restart the Engine to continue the replication process.

4. When no error is reported, the Run dialog opens.
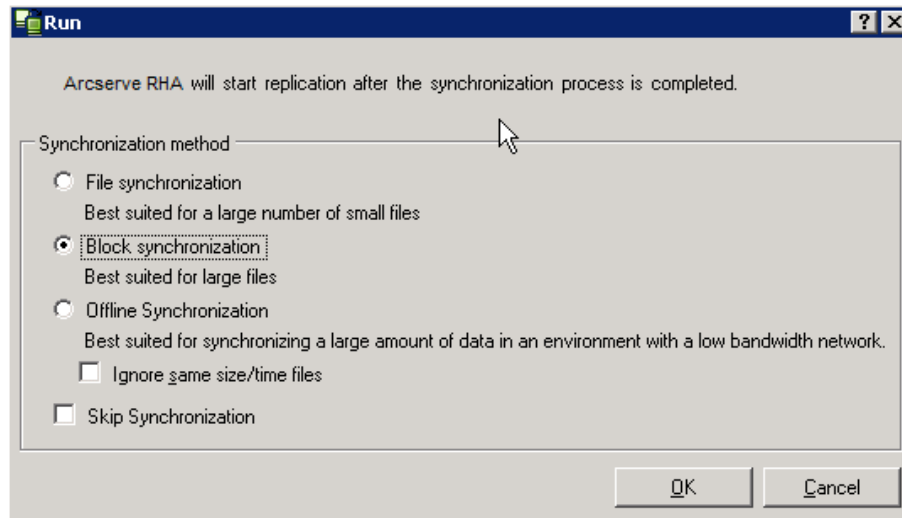


The Run dialog lists the synchronization options. You can choose File, Block or Volume synchronization. For more information, see Synchronization.

**Note**: In general, the default values are the most appropriate choice.

5. For File Server scenarios, verify that the File Synchronization is selected as the synchronization method, and click OK.

**Important!** Do not skip synchronization unless you are absolutely certain that the data in the Master and Replicas root directories is identical.

The Manager now indicates that the scenario is running via the green play symbol to the left of the scenario, and via the scenario state, which turns into Running.



6. After a scenario begins running, a Statistics tab appears at the bottom of the Framework pane, displaying a graphical view of the replication.

7.  By default, after synchronization starts, a Synchronization Report is gen-
    erated. To view the report, refer to Viewing a Report.

    **Note**: You can also generate a Replication Report on a regular basis to monitor
    the replication process on each participating server.

# Run Mode

While replication is running and the running scenario is selected, the Manager screen becomes gray. Statistics are shown in the Framework pane on the right. Events are recorded in the Events pane on the bottom. You may change certain properties while the scenario is running. See Change Scenario Configuration.

**Note**: When working with clusters, you should be aware that Arcserve RHA does not support renaming a Cluster Group while the Engine is running. To ensure the proper operation of clusters with Arcserve RHA, stop the Engine before renaming a Cluster Group.

# Run a Scenario using a Proxy Server

In cases where configuring a Virtual Private Network (VPN) may not be practical, Arcserve RHA allows replication using proxy servers. One example where use of a proxy server may be applicable is if the Master host resides behind a firewall with only HTTP proxy access and the replica host resides on the public internet.

**Important!** A replica server on the public internet has inherent risks which need to be evaluated by each user prior to implementing a scenario using a proxy server.

To configure a scenario that uses a proxy server, enter the proxy details in the Host Connection property of the Master.

**Note:** After you enable proxy settings in the Master, the replica automatically starts using the proxy server. If you do not want a replica to use the proxy server, set the replica Enable Proxy setting to Off.

**Follow these steps:**

1. From the Manager, select the scenario.

2. Select the master host and click the Host Connection property.

3. Change the HTTP proxy Information to On and enter the proxy server details such as the host name, port number, and user credentials.

4. Run the scenario.

# Stop Replication

You must stop a running scenario in order to set or change properties.

**To stop replication**

1. From the Scenario pane, select the scenario you want to stop.

2. To stop the scenario, click Stop  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 becomes Stopped by user, and the Statistics tab is no longer available on the Framework pane.

# Synchronize Master and Replica Servers

The synchronization process can be manually activated at any time, whether replication is running or not. Manual synchronization is recommended in the following situations:

- Before starting replication on servers with large amounts of data and a heavy update rate.

- After a lengthy network failure, if automatic synchronization is not activated.

- After restarting one of the participating servers, if automatic synchronization is not activated.

**To synchronize the Master and the Replica servers**

1. From the Manager, select the scenario you want to synchronize.

2. Click Synchronize on the Standard toolbar, or select the Synchronize option from the Tools menu.

   The Synchronization dialog appears, displaying the synchronization options.

3. Choose the desired options. For more information about the synchronization methods, see Synchronization Methods.

   **Note**: You can also set synchronization to run automatically at pre-scheduled hours on specific days, and exclude it from specific dates. For more information, see Schedule Synchronization.

   Click OK to start the synchronization according to the method you selected.

   **Note**: You can issue a synchronization command while a Replica is suspended; however it is performed only after the replication is resumed.

   After the synchronization process starts, the Statistics tab opens, informing you thatSynchronization is in progress**.**

**Additional Information**:

- Synchronization Methods
- How Offline Synchronization Works
- Schedule Synchronization

# How Offline Synchronization Works

Offline synchronization lets you copy data to be replicated to an external device and then from the device to the Replica server. Offline synchronization is an effective method for transferring huge data volumes with low network bandwidth.

Consider these tips when using offline synchronization:

- You calculate the daily change rate of the data being protected on the Master and multiply that value that is based on the number of days that are required for the data to reach the Replica server. For example, if your daily change rate is 2 GB and it takes 3 days before the data can be applied to the Replica, assume the spool size on the Replica to be about 6 GB. You can run scenarios in Assessment mode to derive this information.

- The spool must have enough space to hold the value you derived in your estimate. If the spool fills, scenarios must be restarted.

- You should choose Continue Offline Sync without verification unless you are certain the copy tool you are using to transfer the snapshot to the replica preserves all security attributes. ACL properties must be identical after copy operations for Offline Synchronization to complete properly. Security attribute mismatches can occur if you use normal copy tools with the verify option. Windows Explorer should not be used as it can change security attributes.

  **Note**: You cannot use offline synchronization for scenarios with scheduled replication or scenarios in which the Master is a UNIX/Linux host.

You can perform offline synchronization for all other types of scenarios. The following examples demonstrate how offline synchronization works for these scenarios.

- Full system, high availability

- Non-full system high availability

- Full System backward or BMR restore

# Perform Full System High Availability Scenarios

Use this method of offline synchronization when the production data set is very large and the WAN link does not have enough throughput to allow for a synchronization of the total data set in a reasonable amount of time.

The offline synchronization process creates VSS snapshots of the volume that resides in the root directory on the Master server. By default, Arcserve RHA mounts the root directory to the location where the RHA Engine was installed, such as the C:\ drive. The following example illustrates the location of the VSS snapshot files when the offline synchronization process runs:

C:\OfflineSyncMountPoint\<scenarioID>\E_volume

```
Synchronization data for 'C:/CA_install_log' is located in
'C:/OfflineSyncMountPoint/1360975302/C_Volume/CA_install_log' and it is ready to be manually
copied to the Replica(s).
```

For the Replica server, you create another File Server DR scenario to replicate from the external device into the target root directory. For example, E:\Data on the Replica server. Optionally, you can perform a Robocopy, (robust file copy), a file copy, or any other copy methods from the Master mounted snapshot volume, and then manually copy the data from the external drive into the target folder on the Replica server after you transferred the drive.

**Follow these steps:**

1. Configure a scenario in the usual manner.

2. When you reach the final stage in the Scenario Creation Wizard, do not click Run Now. Instead, click Finish.

3. From the Manager, click Run. The Run dialog opens.

4. Select Offline Synchronization and click OK.

   After the scenario begins, the Event Console displays the path to the data so you can perform the manual copy. Arcserve RHA takes a VSS snapshot so you can copy data with no impact to the production data set.

   a. The master volume generates VSS and you can find data at C:\OfflineSyncMountPoint\<scenario ID>\.

      **Note:** You can customize the root volume by configuring the SnapshotStorageVolume value in ws_rep.cfg.

   b. On the replica, the disk volume is generated and then mounted as <install path>\Engine\vm\<scenario ID>\<volume ID>\.

c. You should copy all data from master to replica mount point. Manually copy data from the VSS path created on the master to an external device or optical media (DVD-R, CD-R). Physically transport the external media to the Replica server.

During transport, the production dataset caches any changes in the source data in the Replica spool.

5. Copy the data that is stored on the external media to the correct root directory on the Replica server.

Use a copy utility such as xcopy or fastcopy to copy all master server files and folders from the OfflineSyncMountPoint folder to the replica folder <install path>\Engine\vm\<scenario ID>\<Volume>.

xcopy [Source] [Destination] /E /C /H /R

6. After the files are copied, Click Tools, Offline Sync Management to open the Offline Sync Management dialog.

7. Select or clear the "Ignore Files of Same Time/Size" check box and click OK.

8. When the block synchronization process starts, events are shown in the Event Console. When the process is complete, the Synchronization status changes to Replication. Any changes cached in spool on the Replica during the offline sync are now committed to the Replica and deleted from the spool.

# Perform Non-full System High Availability Scenarios

Use this method of offline synchronization when the production data set is very large and the WAN link does not have enough throughput to allow for a synchronization of the total data set in a reasonable amount of time.

The offline synchronization process creates VSS snapshots of the volume that resides in the root directory on the master server. By default, Arcserve RHA mounts the root directory to the location where the RHA Engine was installed, such as the C:\ drive. The following example illustrates the location of the VSS snapshot files when the offline synchronization process runs:

C:\OfflineSyncMountPoint\<scenarioID>\E_volume

Synchronization data for 'C:/CA_install_log' is located in 'C:/OfflineSyncMountPoint/1360975302/C_Volume/CA_install_log' and it is ready to be manually copied to the Replica(s).

**Note:** In the following window, C_Volume represents a VSS mounted volume at this point. For example, when you want to retain the sparse file attributes, you can create a new File Server DR scenario to replicate the data from within this mounted snapshot to an external device.
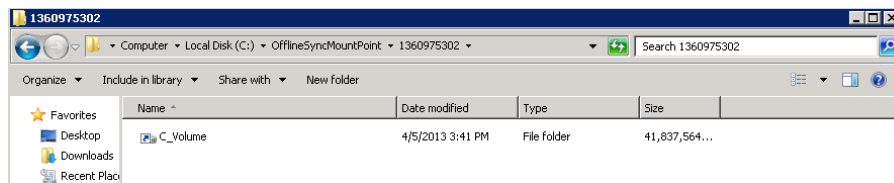


For the Replica server, you create another File Server DR scenario to replicate from the external device into the target root directory. For example, E:\Data on the Replica server. Optionally, you can perform a Robocopy, (robust file copy), a file copy, or any other copy methods from the Master mounted snapshot volume, and then manually copy the data from the external drive into the target folder on the Replica server after you transferred the drive.

**Follow these steps:**

1. Configure the scenario in the usual manner.

2. From the Manager, click Offline Sync.

3. Copy data from the mounted VSS snapshot to external device.

4. Move the device to the replica server and then copy the data to the target folder.

5. From the Manager menu, click Tools and Offline Sync Management

6. Select the offline sync option that you would want to use.

**Note:** The Verification option compares the files in mounted VSS snapshot on the master server with the data that resides in the root directory on the replica server and reports back to you if they are identical or not.

7. Continue running the scenario.

# Perform Full System Backward and BMR Restore Scenarios

Use this method of synchronization when you want to synchronize Bare Metal Recovery (BMR) data at the volume level or block level.

**Follow these steps:**

1. Run the Backward scenario or BMR restore scenario and enter Source and Destination IP address as usual.

2. When you reach the final stage in the Scenario Creation Wizard, select Offline Synchronization.
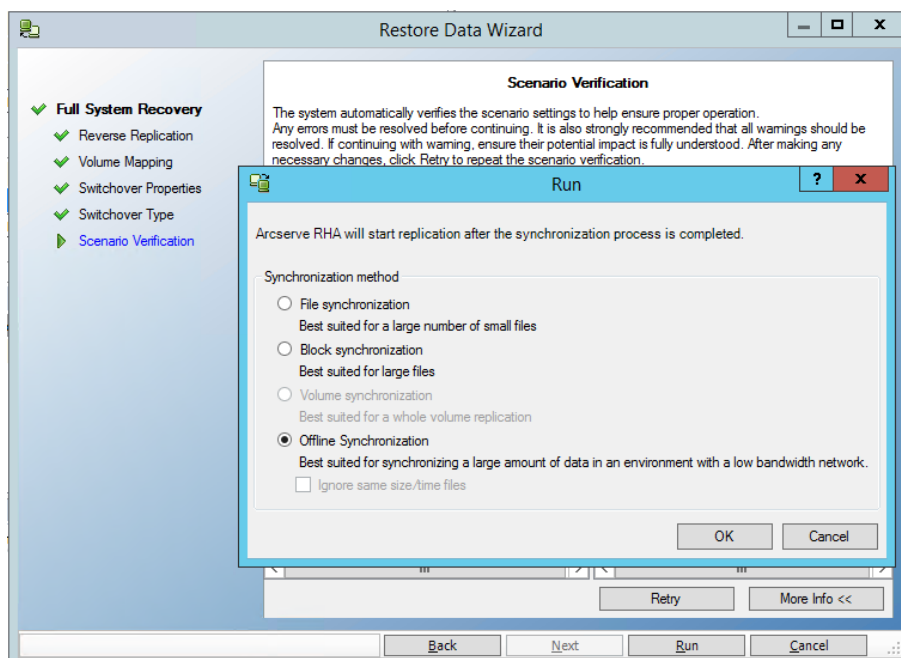


3. After the scenario starts, check the Event log for the folder name.
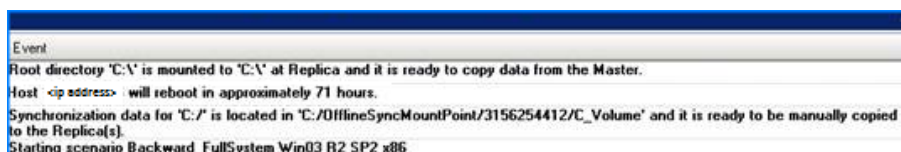


4. Log in to the appliance server and locate the VSS mount point folder.

   For the BMR restore scenario, log in to the appliance server and locate the VSS mount point folder.

   For the backward scenario, log in to the VM generated by switchover and locate the VSS mount point folder.
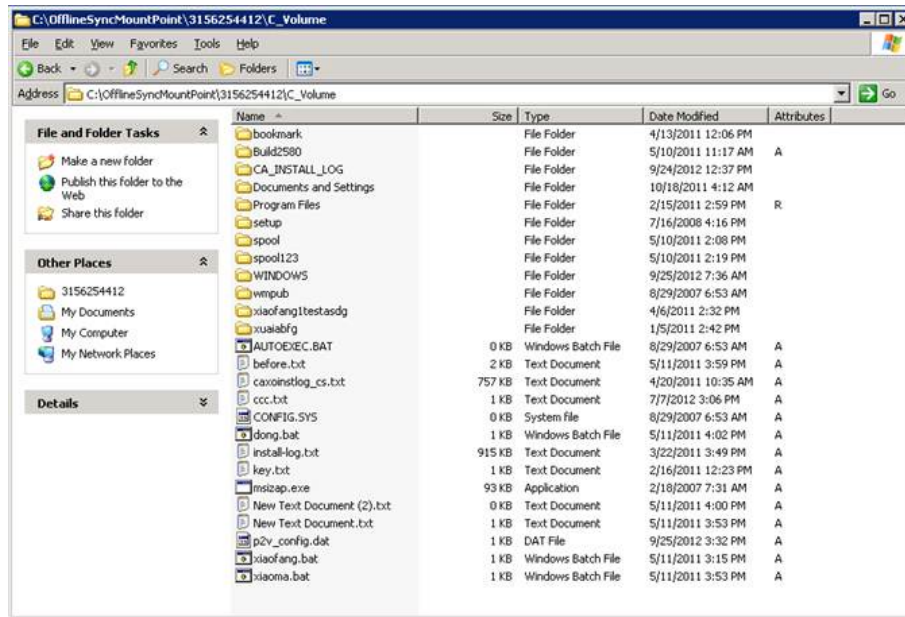
5.  Manually copy data from the VSS path created on the appliance server or the VM generated by switchover to an external device or optical media (DVD-R, CD-R). Physically transport the external media to the BMR server.

    During transport, the production dataset caches any changes in the source data in the Replica spool.

6.  Copy the data that is stored on the external media to the correct root directory on the BMR server mapped volumes.

7.  Use a copy utility such as xcopy or fastcopy to copy all files and folders from the appliance server to the BMR server.

    xcopy [Source] [Destination] /E /C /H /R

    Wait for copying to complete.

8.  Navigate to the RHA Manager and click the Offline Sync Management button.

    The Block synchronization starts.

9.  After the block synchronization is complete, perform the switch process.

# Host Authentication

When you create and run a scenario, Arcserve RHA verifies user credentials of all hosts. Arcserve RHA also ensures that replica files are created in the right folders when multiple users replicate to a single replica server. The authentication process verifies that the user has proper permissions to their folder on the replica server. If the host credentials or folder verification fails, the scenario stops immediately.

This section contains the following topics:

- Enabling Multitenancy Replication
- How to Enable Multitenancy Replication

# Enabling Multitenancy Replication

In a Multitenancy replication setup, multiple masters belonging to unrelated customers are replicated to a single replica server. The replica server belongs to a Managed Service Provider (MSP) and is managed by an MSP administrator. The end user owns a remote Arcserve Control Service (CS). Users create scenarios and replicate data to the replica server.

On the replica server, the administrator creates users, assigns parent folder of the root directory, and grants Full Control permissions to the parent folder. The user name, password, the parent folder details are then shared with users. Users specify these details while creating the scenario.

Users have limited access to the replica machine and have Full Control permissions only to their folder. They can create root directories only in their folder.

The following diagram illustrates a multitenancy setup.



**Multitenancy Setup**
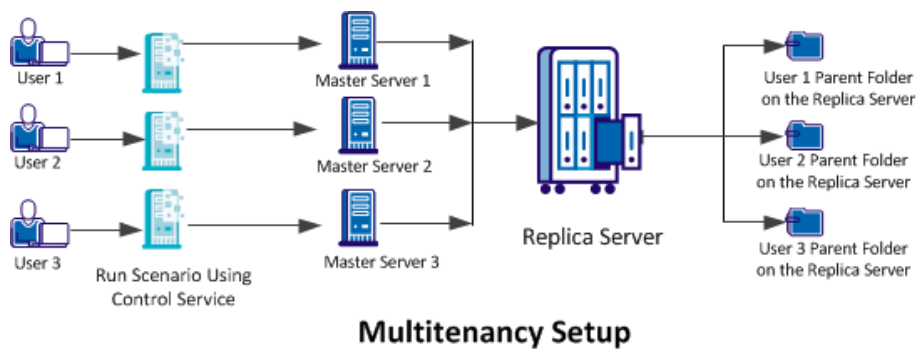
# How to Enable Multitenancy Replication

To enable multitenancy replication, the administrator must first create users in the replica server and grant Full Control permissions to the parent folder of the root directory. Now when a user runs the scenario, Arcserve RHA verifies whether the user has proper host credentials and Full Control permission to the parent folder. The scenario runs successfully if both the criteria are validated otherwise the scenario stops with an error.

The following diagram illustrates how to enable multitenancy replication in an MSP setup.
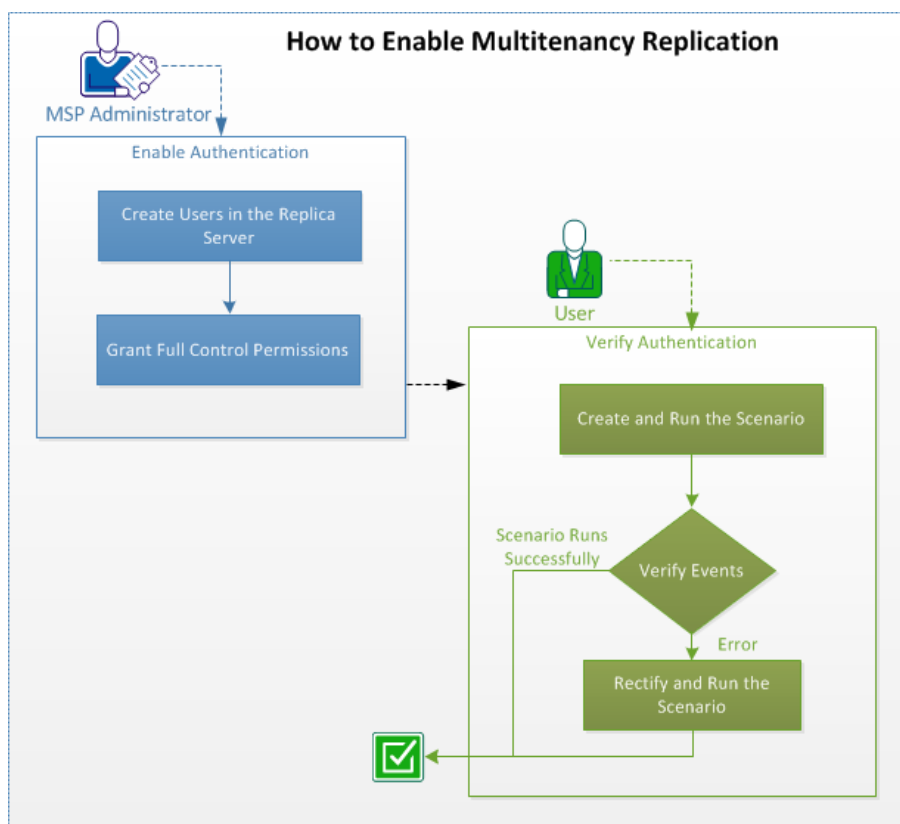


Perform the following tasks to enable multitenancy replication:

This section contains the following topics:

1. Create Users on the Replica Server
2. Grant Full Control Permissions
3. Create and Run the Scenario
4. Verify Events
5. Rectify and Run the Scenario

6. Disable Authentication

7. Considerations and Limitations

# Create Users on the Replica Server (MSP Admin- istrator)

As an administrator, create users on the replica server for each user. You define the parent folder for each user where the user saves the replica root directories which contains replicated files. You can create folder in any location. For example, the following folders can be created as parent folders.

C:\Uploads\User 1

C:\Uploads\User 2

# Grant Full Control Permissions (MSP Administrators)

After you create the folder, make sure that each user stores replica files to their own folders. Grant the Full Control permission to each user on the parent folder of their root directory.

**Follow these steps:**

1. Log in to the Replica server and navigate to the folder where you have created user folders.

2. Right-click a user folder and click Properties.

3. Select the Security tab.

4. Select the user and grant the Full Control permissions.
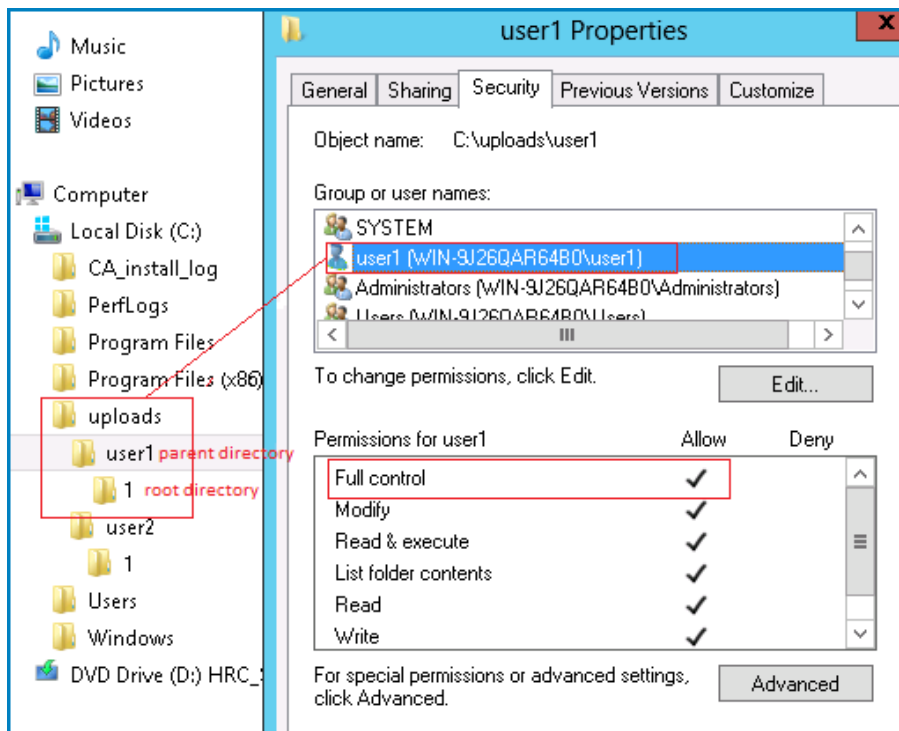


Similarly, select other user folders and grant the Full Control permissions.

# Create and Run the Scenario (Users)

When you create the scenario, enter the valid replica credentials and select the root folder in the specified parent folder. This user credential and folder details are provided to you by the MSP administrator.
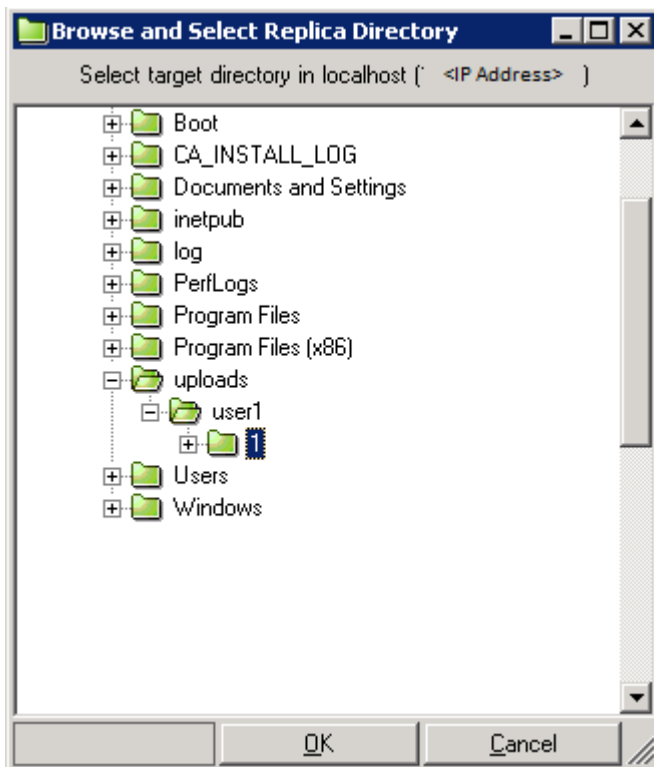


When the scenario runs, RHA verifies the following conditions:

1. Scenario level credentials for each host.

2. Host level credentials for each host. The host level credentials override scenario level credentials. If the host level credential is blank or is not set, then RHA automatically uses scenario level credentials.

3. User has Full Control Permission to the specified parent folder.
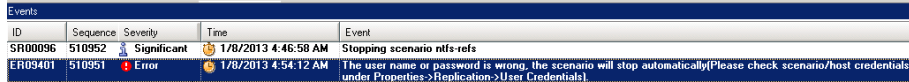
The scenario runs successfully only if the host credentials and folder access permissions are validated. The scenario stops immediately when any of the conditions fail. Rectify the scenario with proper credentials or the valid replica folder.

# Verify Events (Users)

The scenario fails to run when you specify invalid host credentials or you do not have Full Control permissions to the parent folder. Arcserve RHA displays an error when the scenario fails. Verify the error and modify the scenario accordingly.

# Rectify and Run the Scenario (Users)

When the scenario fails to run and Arcserve RHA shows an error, modify the scenario and verify the replica credentials or the valid replica folder.

| Events | | | | |
|---|---|---|---|---|
| ID | Sequence | Severity | Time | Event |
| SR00096 | 510952 | Significant | 1/8/2013 4:46:58 AM | Stopping scenario ntfs-refs |
| ER09401 | 510951 | Error | 1/8/2013 4:54:12 AM | The user name or password is wrong, the scenario will stop automatically(Please check scenario/host credentials under Properties->Replication->User Credentials). |

**Follow these steps:**

1. Navigate to the RHA Manager and select the scenario.

2. From the Scenario view, select the Replica server and click Properties on the right pane.

3. Expand Replication and click User Credentials.

Specify the correct Replica server credentials and select the root directory of the parent folder as provided by the administrator.

# Disable Authentication

By default the EnableAuth parameter in the ws_rep.cfg file is set to True. To disable the host authentication on a specific host, set the EnableAuth parameter in the ws_rep.cfg file to False.

**Follow these steps:**

1. Open the ws_rep.cfg.

2. Modify the EnableAuth parameter to False.

The host authentication is now disabled.

## Considerations and Limitations

Be aware of the following considerations and limitations of a multitenancy replication:

- Local or domain administrators can successfully run scenarios even if they do not have proper permissions on the replica folders.

- While creating a scenario, users can see the directory structure of other users.

- Only the administrators are allowed to run scenario scripts by default. For other users to run the script, set the OnlyAdminCanRunScript in ws_rep.cfg to false.

# Close and Open the Manager during Replication

After the scenario has been defined, and replication has started, the Manager can be closed. It may remain open only for real-time monitoring of the replication process. Closing the Manager does NOT stop the running scenarios. When it is opened again, it automatically uploads all of the saved scenarios and displays their status.

**Note:** Even when the Manager is closed, it is possible to monitor the replicated system via the Overview page. You can also get notifications by email or by automatically running user-defined scripts when important events or errors occur. For more information, see the *Event Notification* topics, in the *Scenario, Master and Replica Properties* sections.

# Suspend Replication

At times, it may be necessary to suspend updates on a Replica machine in order to perform system maintenance or some other form of processing that does not modify the replicated data there. Usually, it is not desirable to stop replication since this requires a resynchronization afterward. The replication suspension feature of Arcserve RHA solves this problem.

During the suspension period, all changes are spooled on the Master or on the Replica located upstream of the suspended Replica. In other words, changes continue to be recorded for update on the suspended Replica, but are not actually transferred until replication is resumed. After replication is resumed, the accumulated changes are transferred and applied without any need to perform a full resynchronization of the data.

Replication may be suspended either manually or on a scheduled basis.

**Important!** It is imperative that during suspension, you do nothing on the Replica that causes the data to change in any way, including starting an application such as Exchange Server, SQL Server, or Oracle. If you need to start programs that will change data on the Replica, you may use the Assured Recovery option.

Be aware of the following:

- You cannot suspend replication during synchronization. You can suspend replication only temporarily, since changes are accumulated in the spool directory of the Master or upstream Replica. Make sure that sufficient disk space is available for the spool to hold the changes during the time the Replica is suspended.

- In a scenario that has more than one Replica host, you can only suspend one Replica at a time.

**More information:**

- Suspend Replication Manually
- Resume Replication after Manual Suspension
- Schedule Automatic Replication Suspension

# Suspend Replication Manually

You can manually suspend replication for system maintenance operations.

**To manually suspend replication**

1. From the Manager, select the Replica you want to suspend. Then, click Suspend, or select the Suspend Replication option from the Tools menu.

   A confirmation message appears, informing you that any change of the Replica root directories content during suspension requires manual resynchronization.

2. Click **Yes** to suspend the replication.

   After the Replica is suspended, a red icon appears next to the Replica on the Scenario pane.

   **Note:** During suspension, the scenario state does not change but stays **Running**, since it is only the replication to the Replica that is suspended.

   On the Scenario Statistics pane, a suspension icon and a caption appear, informing you that replication is suspended.

3. While replication is suspended, you may perform maintenance on the Replica server, including restarting the Replica server. However, it is very important not to modify the replicated data in any way or a resynchronization with the Master will be required.

   **Note**: You can issue a synchronization command while a Replica is suspended; however, it is performed only after the replication is resumed.

# Resume Replication after Manual Suspension

While replication is suspended, changes are spooled on the Master. After you complete system maintenance operations, you must resume replication and end the manual suspension period, which transfers the changes accumulated on the Master to the Replica server.

**To resume replication after manual suspension**

1. After you suspend a Replica, the **Suspend** button on the Manager toolbar toggles to Resume Replication. When you are ready to resume the replication, click the Resume Replication button, or select the Resume Replication option from the Tools menu.

   A confirmation message appears.

2. Click **Yes** to resume replication.

   After replication is resumed, the red icon disappears from the Replica on the Scenario pane, and the suspension symbol disappears from the Scenario Statistics pane.

# Schedule Automatic Replication Suspension

You can schedule replication suspension automatically.

**To schedule automatic replication suspension**

1. From the Manager, select the Replica you want to suspend and stop the scenario in which it participates.

2. On the Framework pane, select the Properties tab to open the Replica Properties list.

3. On the Replica Properties list, open the Scheduled Tasks group. On the Suspension property, set the value to On**.**

4. On the Schedule property, click the value box.

   The Suspend hours dialog opens.

   The Suspend hours dialog is similar to the Schedule Setting dialog, which is used for scheduling automatic synchronization. For information about setting a schedule, refer to Schedule Synchronization.

5. Set the schedule for automatic suspension in the Suspend hours dialog, and click OK to save your schedule and close the dialog.

6. To activate the schedule, click the Save button on the Standard tool bar and start the scenario.

   The Replica you selected for suspension will be suspended and resumed according to the set schedule.

# Running a Scenario in Assessment Mode

Arcserve RHA enables you to run a scenario without actually replicating data, in order to assess the accurate bandwidth usage and compression ratio benchmarking that is needed for replication. When you run a scenario in Assessment mode, no replication occurs, but statistics are gathered. A report is provided after you stop the assessment process.

To use the Assessment mode option, you need to create a new scenario, and select the **Assessment Mode** check box in the Scenario Creation Wizard.

A scenario that can run in Assessment mode can also run in a regular replication mode. When initiating the scenario, the button you click, either the green **Run** button or the blue **Run (Assessment mode)** button, determines the running mode of the scenario.

**To run a scenario in Assessment Mode**

1. From the Manager, select the scenario you want to run in Assessment mde.

2. To run the scenario, click **Run (Assessment mode)** on the Standard toolbar.

   Arcserve RHA verifies the scenario before running it.

3. If the scenario was not set up correctly or problems occurred in the participating hosts, errors are reported on the Event pane. When no error is reported, the scenario starts running.

   The Manager now indicates that the scenario is running in Assessment mode via the blue play symbol to the left of the scenario, and via the scenario's state which turns into **Simulation**.

4. After the scenario begins running, a Statistics tab appears at the bottom of the Framework pane, displaying a graphical view of the replication assessment.

5. By default, after you stop a scenario that runs in Assessment mode, an Assessment Mode Report is generated. To open the report, refer to Viewing a Report.

# Chapter 6: Monitoring Replication

This section describes the various monitoring tools of Arcserve RHA that enable you to control and monitor your replication environment.

This section contains the following topics:

# The Overview Page

The Overview Page allows simultaneous monitoring by any number of administrators or professionals who wish to be informed about the state of the replicated system.

**Note:** There is no need to refresh the browser; the Overview page refreshes automatically.

# The Manager

The Manager lets you control and monitor your replicated system.

This section contains the following topics:

- [How to Monitor Multiple Scenarios](#)
- [State Information](#)
- [Live Statistics](#)

# How to Monitor Multiple Scenarios

The Scenario pane shows the current status for all scenarios in one pane.

You can customize the way the columns are displayed. For more information, see Customize the Scenario View.

# State Information

State information is displayed beside each scenario name and beside each server in the replication tree, whenever a synchronization process has been started or completed, and whenever a replication process is under way.

The state information includes:

- A graphic indication next to the scenario name indicating the scenario is running , or is idle .

- A graphic indication next to the server name indicating that the server is a Master (active) server , or a Replica (standby) server .

- A graphic indication of whether the servers are connected: if the connection to any of the participating servers is lost, the server icon appears with a large red **X** marked over it.

# Live Statistics

Live statistics are displayed in two areas:

- Scenario pane
- Statistics pane

This section contains the following topics:

- Scenario Pane
- Statistics Tab
- Refresh Statistics Display Automatically
- Refresh Statistics Display Manually

# Scenario Pane

The live statistics displayed in the Scenario pane include the following information:

- **Changed** - total data replicated from this host since the last synchronization.

- **Synchronized** - total data synchronized from this host.

- **Files** - total number of files replicated from this host.

- **In Spool** - total (current) amount of data contained in the spool.

# Statistics Tab

The Statistics tab in the Framework pane also displays live statistics. Different statistics information is displayed for a scenario, a Master and each Replica server.

**Note**: The Statistics tab on the Framework pane appears only when a scenario is running.

The available live information in the Statistics tab is as follows:

▪ **Statistics per scenario** - a graphical overview of the scenario state.

▪ **Statistics per Master** - a table that containing the following information: state; replication starting date and time; Arcserve RHA Engine version no.; total amount of data contained in the spool; total amount of data changed in the root directories since the beginning of the replication process, including the number of created folders and changed, removed and renamed files; total size of files sent and replicated from this host; synchronization method; and synchronization progress.

When running synchronization for very large files, additional tables appear, displaying in details the synchronization progress for each file in each root directory.

Be aware of the following:

◆ The definition of a large file depends on the value of the BDMaxFileSizeToSendWholly property. This property is stored on the Engine machine, in the INSTALLDIR\Engine\ws_rep.cfg file. The default is 10MB. When a synchronized file is smaller than this value, it will not appear on the table.

◆ The appearance of the detailed Synchronization Progress table also depends on the value of the UseNewSynchStatistics property.This property determines whether the detailed Synchronization Progress table will be displayed when there are large files. When the value of this property is True, the table will be displayed. The default value is True, and the property is also stored in the ws_rep.cfg file.

The Synchronization Progress tables contain the following information for each synchronized file: synchronization state; file name; file size; amount and percentage of data that was already compared between the Master and Replica; amount of data that needs to be sent from the Master to the Replica; amount and percentage of data that is the same on the Master and Replica, and therefore is not sent to the Replica; synchronization starting date and time; summary of the synchronization progress of each root directory.

**Note**: When the File Synchronization method is running, the files are synchronized in their entirety. Therefore, the Already Checked column can contain only two values - 0% or 100% - for an individual file.

Each synchronized root directory is represented by a separate Synchronization Progress table, and each Synchronization Progress table can display statistics of 10 files at the most. When a root directory contains more than 10 files, the 10 largest files will be presented in the table.

▪ **Statistics per Replica** - a table that containing the following information: state; replication starting date; Arcserve RHA version no.; total amount of data contained in the spool; and total amount of data (in KB) changed in the root directories since the beginning of the replication process, including the number of files changed, removed and renamed.

# Refresh Statistics Display Automatically

The Manager receives state information from all servers participating in the current scenario. You can set a default frequency for automatically updating the state information and live statistics display.

**To define the refresh rate**

1. From the Tools menu, select Statistics, Refresh Rate.

   The Refresh Rate dialog opens.

2. Enter the desired refresh rate in seconds and click OK. The Scenario pane updates accordingly.

   **Note**: The refresh rate interval can be between 10 to 99 seconds.

# Refresh Statistics Display Manually

**To manually refresh the displayed information**

1. Click Refresh Statistics on the Standard toolbar on the Manager screen.

2. Press the F5 key.

3. From the Tools menu, select Statistics, Refresh.

   The statistics data updates.

# View Events

The Events pane displays messages and general information about the selected scenario. This information is received from the servers participating in the running scenario. The information displayed can inform you, for example, that a directory is synchronized, a server is connected, synchronization started/finished, and so on. The information includes the server name and time, and a brief explanation. Important events or error messages are shown in bold letters.

**More information:**

- View Events in a Separate Window
- View Incoming Events
- Copy Events for Use with Other Programs
- Filter Events

# View Events in a Separate Window

Sometimes, event messages are very long, exceed the Event area, and are cut off (visually). In these cases, you may want to display the event messages in separate windows.

**To view events in a separate window**

1. From the Manager Event pane, select the event you want to view.

2. Double-click the event, or right-click it and select View Event in other Window from the pop-up menu. Alternatively, select from the Events menu the View Event in other Window option.

   A pop-up message appears, displaying the full message text of the selected event.

# View Incoming Events

The Manager can visually notify you when an incoming event occurs.

**To view incoming events as they occur**

1. From the Manager Events menu, select the Pop-up on Incoming Event option.

   When an incoming event occurs, the Manager icon in the task bar flashes and the Manager is moved to the foreground.

   **Note**: If you re-select this option and turn it off, the minimized application in the task bar does not flash during an incoming event.

# Copy Events for Use with Other Programs

When important events occur, you may want to copy their messages to other programs.

**To copy events for use in other programs**

1. On the Event pane, select any number of events using the Ctrl key.

2. Right-click in the Event pane and select Copy, or select the Copy option from the Events menu. Alternatively, press Ctrl+C.

   You can paste the copied event texts into any program you wish.

   **Note**: Arcserve RHA also enables you to copy the event message directly to a file with CSV extension, such as Excel files. After you select the Copy To CSV option, the application that is defined in your computer as CSV opens, displaying the copied message. (To set your default CSV application, see My Computer, Tools, Folder Options, File Types, CSV.)

# Filter Events

You can filter the events that will be displayed in the Event pane according to different criteria.

**To filter displayed events**

1. On the Event pane, right-click and select Event Filters from the pop-up menu, or select the Event Filters option from the Events menu.

   The Event Filters dialog opens.

2. Use one of the following criteria to filter the events that will be displayed in the Event pane:

   - **Severity** - clear the severity level check boxes that you do not want to display, or select the severity level check boxes that you want to display.

   - **Date** - select the Exclude events older than check box, and then select the time unit (hours/days/months) and the number of units.

   - **Text** - in the A word or phrase in the event box, enter the word or the phrase that you want the displayed events to contain. You can use an asterisk (*) to select any number of characters/digits of all types.

3. To apply the criteria you selected and close the dialog, click OK.

   Only the events that meet the criteria you defined are now displayed in the Event pane.

4. To clear the existing criteria and display all events, on the Event Filters dialog click Reset, and then OK.

# Arcserve RHA Reports

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. To set these options, see the following*:*

- For defining the storage directory and the retention period of the reports, see Understanding Scenario Properties.

- For defining the automatic generation of synchronization and replication reports for the Master, see Understanding Master Properties.

- For defining the automatic generation of replication reports for the Replica, see Understanding Replica Properties.

By default, Arcserve RHA stores reports in the following directory: *[ProgramFilesFolder]\CA\Arcserve RHA\Manager\reports*

**Important:** A report cannot be created if the scenario name contains special characters (i.e. \/?:"<>|,).

This section contains the following topics:

- View a Report
- Deleting Reports
- Synchronization Reports
- Replication Reports
- Open a Backward Scenario Report
- Create Difference Reports
- Assessment Mode Reports

# View a Report

The various Arcserve RHA reports let you determine scenario status and manage your environment.

**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.

    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.

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.

    **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 opens.

# Deleting Reports

The reports that are displayed in the Report Center are stored for the period that is defined in their scenario properties, under the Report Handling group. The reports are displayed even if their scenarios were removed from the Manager. However, the Report Center enables you to delete reports that are no longer of use to you.

**To delete a report**

1. On the Report Center, select from the Available Reports per Scenario table the scenario whose reports you want to delete.

2. Click Delete at the right end of the line.

   A confirmation message appears, asking you whether you want to delete the row.

3. Click OK on the confirmation message.

   The selected scenario is removed from the Report Center, and all of the reports that belong to it are deleted.

# Synchronization Reports

Following synchronization, Arcserve RHA creates a report that lists files that have been transferred or modified. The first few lines (at the top) include: the synchronization method, the scenario name, the names of the Master and the Replica (s), and the synchronization date.

The Summarized Synchronization Report shows the sum total of the removed and modified files as well as bytes transferred. The report also provides information about the number of new directories that were created on the Replica following synchronization, and the number of synchronized files that have different Windows security properties.

The Detailed Synchronization Report presents the complete listing of the files that were transferred or modified during the synchronization process. For each file, the following information is provided:

- **Event** - the action that was performed on the Replica.
- **Bytes** - the size of the file.
- **Time Stamp** - modification time.
- **File Name** - the name and full path of the file.

# Replication Reports

The Replication Reports are generated periodically, according to a pre-defined frequency, and they are generated individually for the Master and the Replica servers.

The Replication Reports include statistics on data replicated since the beginning of the replication process, as well as statistics on data replicated since the last report. The data includes the number of replicated bytes, and the number of files created/updated/removed/and renamed. You can view either a summarized or a detailed report.

By default, Replication Reports are NOT automatically generated. To schedule Replication Reports, see the topic, Report Properties.

# Open a Backward Scenario Report

A Backward Scenario Report is actually a Synchronization Report, which displays synchronization information about a Backward scenario.

**To open a Backward Scenario Report**

1. On the Report Center, select the Backward scenario from the Available Reports per Scenario table.

2. On the Reports table, select the Synchronization Report you want to view.

3. The Synchronization Report for the Backward scenario opens.

# Create Difference Reports

A Difference Report compares the differences between the Master and the Replica at a certain point in time. The comparison is performed using the same algorithms that are used in the synchronization process, but no data is transferred. A Difference Report is generated for each Replica and sent to the Manager at the end of the process. This can be produced at any time.

**Note**: When you run a difference report on a SQL Server 2008/Windows Server 2008 R2 platform, differences are reported due to way in which SQL Server 2008 pre-allocates buffer size. This difference does not impact database integrity.

**Important!** We do not recommend initiating the Difference Report when data is being updated on the Master, since all updates that are not yet applied to the Replica will be shown as difference.

**To create difference reports**

1. Click the Difference Report icon on the Standard toolbar, or select from the Tools menu the Report, Difference report option.

   The Difference report dialog opens with the same comparison methods as in the Synchronization method dialog.

2. Choose the desired options, according to the instructions specified in How Synchronization Works.

3. Click OK.

   At the end of the process, a Difference Report is generated for each replica, and all the reports are sent to the Report Center.

# Assessment Mode Reports

Arcserve RHA generates Assessment Mode Reports after you stop a scenario that runs in Assessment mode. This Report displays statistics about the sum total of bytes that were supposed to transfer from the Master to the Replica, from the initiation of the scenario until it is stopped.

The sum total of bytes that were supposed to transfer is calculated for pre-defined time units, called "Time Stamps". By default, a Time Stamp is calculated every 10 minutes.

**Note**: You can change the Time Stamp default interval in the **ws_rep.cfg** file, by entering a new value for the **AssessmentTimeSample** parameter.

The **Summary** table of the Assessment Mode Report shows statistics about the maximum, minimum and average size of the data that was supposed to transfer during the scenario run. It also provides statistics about data that was supposed to transfer in compressed form.

# Chapter 7: Editing and Managing Scenarios and Hosts

This section demonstrates the manual configuration of a generic File Server replication scenario, and explains the auto-discovery process for database applications. For more detailed instructions scenarios tailored to specific applications such as Exchange Server or SQL Server, see the appropriate Operation Guide.

After you create a scenario using the Scenario Creation Wizard, you can manually edit most of its definitions. Although you cannot manually create a scenario from scratch, you can click the **Finish** button at any point, close the wizard and continue the configuration manually.

This section contains the following topics:

# Define the Master and Replica Servers

Every scenario is automatically created with one Master and one Replica.

**To define the Master or Replica host**

1. On the Scenario pane, right-click the **Enter Master/Replica host name here** text and select **Rename** from the pop-up menu. Alternatively, double-click the required text.

   Enter the hostname or IP address of the host.

2. Press the **Enter** key, or click anywhere outside of the text field.

3. Save your changes by clicking the **Save** button.

   After defining a new host, you need to define its root directories for the data replication.

   ◆ To define the Master root directories, refer to Select Master Directories and Their Contents for Replication.

   ◆ To define the Replica root directories, refer to Select Replica Root Directories.

# Add Additional Replica Servers

When creating a scenario using the Scenario Creation Wizard, you can define only one Replica server for the scenario. To add more Replica servers to the scenario, follow the instructions below.

**To add additional Replica servers**

1. On the Scenario pane, select the host (Master or Replica) under which you want to add a server. Right-click it and select**Insert Host** from the pop-up menu, or select **Insert Host** from the **Edit** menu.

   A new Replica server entry opens.

2. Define the new Replica as you defined the other Replica host, and set its properties and root directories.

3. Save your changes by clicking the **Save** button.

# Select Master Directories and Their Contents for Replication

This section explains how to select directories and files on the Master for replication.

**Notes:**

- Working with **Master Root Directories** dialog is possible only if the Engine is installed and running on the host.

- You can also select registry keys for synchronization, as described in Synchronize Registry Keys.

**Important**! Special limitations apply to UNC paths (\\server\share) of remote root directories. This path type is not supported as a source (on the Master) for real-time replication. However, it can be the target for data replicated in real-time, meaning it can be used to store data on the Replica. In this case, these root directories can even support ACL replication.

**To select Master directories and their contents**

1. In the Scenario pane, select the name of the Master server whose data you want to replicate.

2. In the Properties pane, click the **Root Directories** tab at the bottom. The Master Root Directories information opens in the pane.

3. Right-click anywhere in the pane, and select **Browse and Select Directories**. Alternatively, double-click the Master root directory named **Directories**.

   The **Master Root Directories** dialog opens.

   The **Master Root Directories** dialog has two areas. The left area shows only directories and sub-directories. The right area shows both directories and sub-directories, and files in those directories. The checkboxes are for you to select or clear. When selected, those directories or files will be replicated. Those not selected are ignored.

4. In the dialog's left area, select the directories that are to participate in the Master's replication scenario, by clicking on the relevant checkboxes. These are the Master root directories. The checkbox is selected, and the directory name is bold:

   **Notes**:

   - When you select root directories for the Master or Replica servers, the total character length of root directory plus subdirectory names should

not exceed 1024 bytes.

- ◆ If a root directory is a sub-directory, it remains bold and marked, and its parent directory is marked with a grayed checkmark.

All files and sub-directories belonging to the directory that is highlighted in the left area are displayed in the right area.

5. You can clear the check boxes of the sub-directories and specific files that appear on the right area. They are then ignored from replication.

**Note**: If you clear any of the right area sub-directories and files, they are ignored, but the root directory is still selected. However, it is marked with a grayed checkmark.

6. When you have finished selecting all the directories and files you want to replicate, click **OK**.

The selected directories now appear in the Root Directories pane under the Master root directories column.

**Note:** When working with SQL Server replication, databases that are added to the already-selected root directories after replication starts, will not be replicated.

# Edit Directory Names

You can edit the names of the Master root directories. However, when changing a root directory name, you need to verify that such a directory actually exists on the Master before you run the scenario. If you try to run a scenario with a non-existent Master root directory, the scenario will not run and a critical error will be reported.

**To edit a directory name**

- On the Root Directories tab, select the directory and enter a new name using Windows conventions;

  - or -

- Right-click the directory name, and select **Rename** from the pop-up menu.

# Remove Master Root Directories

**To remove a Master root directory**

- Right-click a directory entry and select **Remove Directory** from the pop-up menu.

# Filter Master Directory Files

The filtering options enable you to include or exclude files from the Master root directories. These options do not select (or clear) items in the **Master Root Directories** dialog. That has to be done manually. However, the filtering options enable you to fine-tune your directory selection and display only the files that will be replicated.

For example, if you choose to include only text files, you need to select the required directories and enter the filter parameter. Then, only text files that are stored in these directories will be displayed on the left area of the **Master Root Directories** dialog.

The Master root directories filters enables you to use a variety of filtering characters, such as characters, strings, wildcards, file names or extensions, etc. The following standard wildcards are available:

**Note**: A "character" in this context refers only to alphabetical or numerical character.

- An asterisk (*) selects any number of characters/digits of all types.
- A question mark (?) selects any single character or numeric digit.
- A pound sign (#) selects itself or any numeric digit.
- An "at" sign (@) selects itself or any single alphabetic character.
- Entering other characters (one or many) selects for those specific characters.

A given filter selection applies to all files in all selected directories in the scenario.

The Filter options are as follows:

- **No filter** - all the directories and files you manually select will be replicated. This is the default option. Refer to Select Master Directories and Their Contents for Replication.
- **Include files** - ONLY the selected files or file-types will be replicated. Refer to Include Files.
- **Exclude files** - ONLY the selected files or file-types will be excluded from replication, and all others will be included. Refer to Exclude Files.

# Include Files

When using **Include files**, only the files or file-types entered into the Filter box are included in the replication scenario, and only if they are selected (checked). You need to manually select the directories in which these file are stored, and if you manually clear a file check box, you override the **Include files** option.

**To include files**

1. On the **Master Root Directories** dialog, manually select the directories on which you want to apply the filter.

   **Note**: Alternatively, you can manually select the directory check box AFTER you enter the filter parameters.

2. Click the **Include files** option button at the top of the **Master Root Directories** dialog. The Filter box is enabled with an asterisk (*) wildcard.

3. Enter the file types you want to include in the Filter box using the appropriate filtering characters. For example, include all files ending with the extensions *.txt *.exe. Separate the extensions using a space.

   **Note**: Do NOT use a comma or a semi-colon to separate extensions. If a file name includes blanks, enclose the complete file name between quotation marks ("").

4. Click the **Apply** button to filter the directories you selected according to the filter parameters.

   The only files that are displayed on the right area are those that meet the filtering criteria.

5. [Optional] You can manually select or clear individual directories and files. This action overrides the **Include files** option regarding the individual directory or file.

6. Click **OK** to save your root directory selection and close the **Master Root Directories** dialog.

# Exclude Files

When using **Exclude files,** the replication scenario ignores any files that have been filtered out (excluded), and includes all others.

**To exclude files**

1.  On the **Master Root Directories** dialog, manually select the directories on which you want to apply the filter.

    **Note**: Alternatively, you can manually select the directories AFTER you enter the filter parameters.

2.  Click the **Exclude files** option button at the top of the **Master Root Directories** dialog. The Filter box is enabled with an asterisk (*) wildcard.

    Enter the file types you want to exclude in the Filter box using the appropriate filtering characters. For example, exclude all files ending with the extensions *.bat *.dll. Separate the extensions using a space.

    **Note**: Do NOT use a comma or semi-colon to separate extensions. If a file name includes blanks, enclose the complete file name between quotation marks ("").

3.  Click the **Apply** button to filter the directories you selected according to the filter parameters.

    The excluded files are not displayed on the right area, and the displayed files are the ones that will be replicated.

4.  [Optional] You can manually select or clear individual directories and files. This action overrides the **Include files** option regarding the individual directory or file.

5.  Click **OK** to save your root directory selection and close the **Master Root Directories** dialog.

# Synchronize Registry Keys

In addition to synchronizing and replicating application data, Arcserve RHA enables you to synchronize the Master and the Replica registry keys. Using the Registry Synchronization option, you can select which registry keys on the Master will be copied to the Replica, and define the synchronization frequency. You can copy the Master registry keys to the same location on the Replica, or you can change the name and storage path of the synchronized keys. If there are multiple Replica hosts in the replication tree, the registry synchronization process is applied to all of them. The registry keys are not replicated in real time. They are copied from the Master to the Replica on a scheduled basis, according to the frequency you defined.

**Important!** Use this feature with caution. Changing registry keys may result in system failure.

**Notes**:

- This feature does not apply to applications that block the access to their registry keys, or to applications whose registry keys cannot be altered.

- By default, the Registry Synchronization option is disabled.

There are several steps in configuring and running the Registry Synchronization option:

1. Activating the Registry Synchronization property.

2. On the Master host, selecting the registry keys to be synchronized.

3. [Optional] On the Replica host, selecting the name and storage location for the synchronized registry keys.

4. Running the scenario to start the registry keys synchronization.

# Activate the Registry Synchronization Option

The first step in configuring and running the Registry Synchronization option is activating this option and defining its frequency.

**Note:** To configure Registry Synchronization properties, the scenario must be stopped. To run scenarios that include Registry Synchronization, you must run Arcserve RHA using a domain administration account.

**To activate the Registry Synchronization property**

1.  In the Scenario pane, select the name of the scenario for which you want to activate the **Registry Synchronization** property.

2.  In the Properties pane, click the **Properties** tab at the bottom.

    The **Scenario Properties** list appears in the pane.

3.  Open the **Replication** group, select the **Registry Synchronization** property, and set its value to On.

    The **Synchronization Frequency** property appears under the **Registry Synchronization** property.
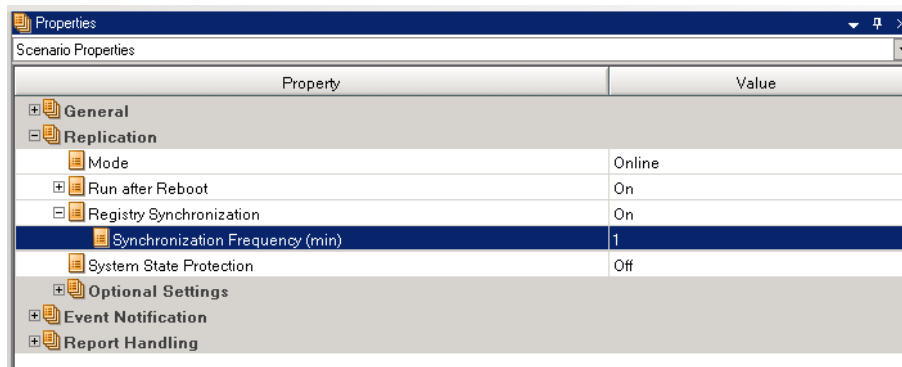


4.  In the **Synchronization Frequency** value box, enter the number of minutes that will pass between each registry keys synchronization.

5.  Save your configuration by clicking the **Save** button on the Standard toolbar.

    Now, you need to <u>select on the Master host the registry keys that will be synchronized</u>.

# Select Registry Keys for Synchronization

The second step in configuring and running the Registry Synchronization option is selecting the registry keys on the Master host that you want to synchronize.

**Notes**:

- To select registry keys for synchronization, the scenario must be stopped.

- You cannot select registry keys for synchronization through the Scenario Creation wizard, only from the Properties pane of the Manager.

- Only keys are displayed for selection. You cannot select specific values for synchronization.

**To select registry keys for synchronization on the Master**

1. In the Scenario pane, select the name of the Master host whose registry keys you want to synchronize.

2. In the Properties pane, click the **Root Directories** tab at the bottom. The **Master Root Directories** information appears in the pane.

3. Right-click the registry object that corresponds to your operating system and select **Browse and Select Registry Keys** from the shortcut menu. Alternatively, double-click the **Registry** object that corresponds to your operation system - either **x86** or **x64**

   The **Browse and select registry keys** dialog opens and displays the Master host registry keys list.
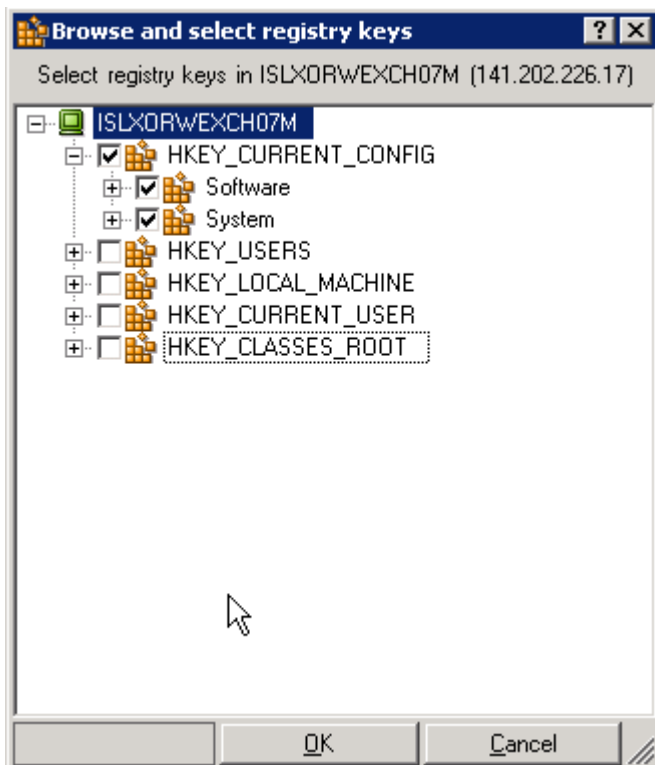
4. Click the checkboxes of the registry keys you want to synchronize.

   **Notes**:

   ◆ You cannot filter registry key selection.

   ◆ If you are manually entering a name and a path of a registry key that
     does not exist on the Master, the scenario verification may be successful,
     but the scenario stops running and an error message is issued. You should
     only enter the details of existing registry keys for replication.

5. When you have finished selecting all the registry keys you want to syn-
   chronize, click **OK**.

   The selected registry keys now appear in the Root Directories pane under the
   **Master Root Directories** column.

6. Save your configuration by clicking the **Save** button on the Standard toolbar.

   By default, the system automatically configures the Replica registry keys to be
   the same as the selected Master registry keys. If you want to change the name
   and storage location of the synchronized Replica registry keys, follow the
   instructions described in the next section.

# Select the Name and Storage Location of the Synchronized Registry Keys

The third step in configuring and running the Registry Synchronization option is selecting the name and storage location of the synchronized registry keys on the Replica host. The system configures the Replica registry keys to be the same as the selected Master registry keys by default, so this step is optional.

**To select the path for storing the synchronized registry keys**

1. In the Scenario pane, select the name of the Replica where you want to store the synchronized registry keys.

2. In the Properties pane, click the **Root Directories** tab at the bottom. The Replica Root Directories information appears in the pane.

   The registry keys selected on the Master for synchronization appear on the Replica in the same location and under the same name.

3. You can change the default path and name of the Replica registry keys in two ways:

   ■ Replacing the default path and name with the path and name of existing registry keys:

      ◆ Right-click anywhere in the pane, and select **Browse and Select Registry Keys**. Alternatively, double-click the name of the specified Replica register key.

         The **Browse and select registry keys** dialog appears.

      ◆ Select the checkboxes of the registry keys you want use, and click **OK** to save your selection.

   ■ Manually entering new path and name for the default values: double-click the registry key name on the Root Directories pane, and manually enter a new path and name.

4. Click the **Save** button on the Standard toolbar.

   To start the registry keys synchronization, you need to run the scenario.

# Auto-discover Database Files for all Databases

To facilitate easy directory selection for standard databases that are supported by Arcserve RHA, database directories and files are identified in the scenario by using database APIs. Arcserve RHA displays the structure of the database and makes the appropriate selections, which you may modify, if necessary.

The Arcserve RHA auto-discovery function automatically discovers all database objects, related files and directories on your database or mail server - whether local or on a network. This function is currently available for all supported applications.

**Note:** Auto-discovery is possible only if both the Arcserve RHA Engine and the database are installed and running on the Master server.

**To use auto-discovery for selecting database files**

1.  On the scenario pane, select the scenario whose database you want to auto-discover, and make sure it is NOT running.

2.  On the Framework pane, open the **Root Directories** tab for the Master.

    The **Auto-discovered** database files icon appears in the Master Root Directories tab.

3.  To start auto-discovery, double-click the **Auto-discovered** icon. Alternatively, select the **Auto-discovery of Database Files** option from the **Edit** menu.

    **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.

    The **Autodiscovery** dialog opens.

    The **Auto-discovery** dialog displays all database directories and files that were auto-discovered.

4.  Select the checkboxes beside the items you want to replicate, and clear the items you want to exclude from replication. Then, click **OK** to save your root directory selection and close the **Auto-discovery** dialog.

# Select Replica Root Directories

You must select the Master directories before this function becomes available. For each Master root directory, you must define a Replica root directory on each related Replica.

**Important!** Special limitations apply to UNC paths (\\server\share) of remote root directories. This path type is not supported as a source (on the Master) for real-time replication. However, it can be the target for data replicated in real-time, meaning it can be used to store data on the Replica. In this case, these root directories can even support ACL replication.

**Note**: Browsing for a directory is possible only if the Engine is installed and running on the selected server.

**To select Replica root directories**

1. In the Scenario pane, select the name of the Replica where you want to store replicated data.

2. In the Properties pane, click the **Root Directories** tab at the bottom. The Replica Root Directories information appears in the pane.

   **Important!** The Scenario Creation Wizard automatically configures the Replica root directories to be the same as the Master root directories. If you want to keep this configuration, ensure that your Replica server has the same drive letters as the Master server, and that the selected directories on the Replica do not contain data you want to save.

3. To change the default Replica root directories, right-click anywhere in the pane, and select **Browse and Select Directories**. Alternatively, double-click the name of the specified Replica root directory.

   **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 connects to the Replica server and opens the **Browse and Select Replica Directory** dialog opens.

   The **Browse and Select Replica Directory** dialog displays the Replica server's directory list.

4. Select a Replica directory to hold the corresponding Master directory. Repeat this for each Master directory.

5.  In order for the replication process to succeed, verify that the user under which the Engine is running has permission for each replication root directory.

    **Note:**The Replica root directory does not have to actually exist. You can enter the directory name by selecting the entry using the standard Windows conventions and Arcserve RHA creates it when the replication starts.

6.  Click **OK** to save your selection and close the **Browse and Select Replica Directory**.

# Propagating Master Root Directories to Multiple Replica Hosts

Arcserve RHA enables you to propagate the root directories you set for the Master to multiple Replica hosts at once. Instead of separately configuring the root directories of each Replica host, in a click-of-a-button you can distribute the root directory of one Master to as many Replica hosts as you want. This option is particularly useful for a scenario that has many Replica hosts.

**Note**: The scenario must be stopped to apply root directory changes.

**To propagate root directories**

1.  On the Scenario pane, select the Master whose root directories you want to propagate.

2.  On the Framework pane, click the **Root Directories** tab at the bottom. The Master Root Directories information appears in the pane.

3.  On the Master Root Directories pane, right-click the root directory you want to propagate and select **Propagate Value** from the shortcut menu.

4.  Click the **Propagate Value** command. The **Propagate Value** dialog opens.
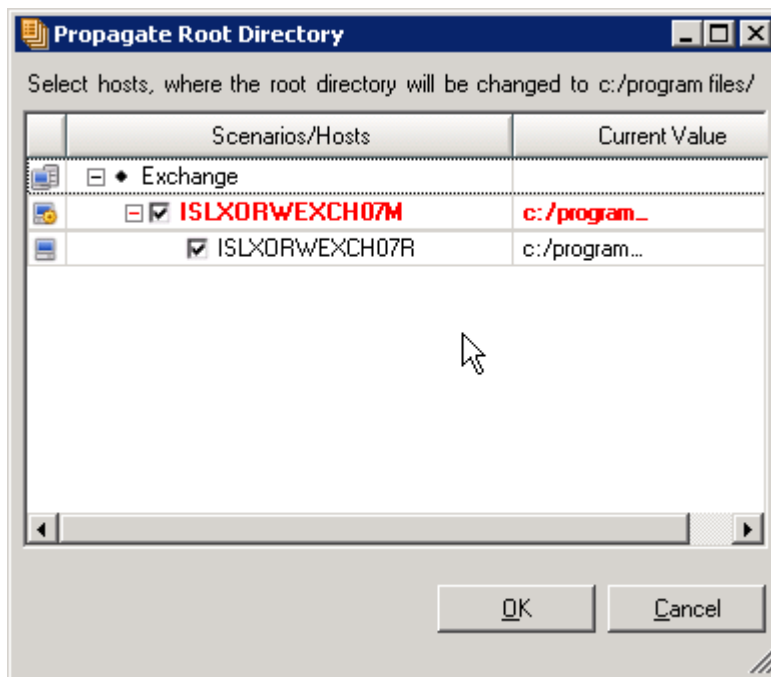


The Master and all Replica hosts in the scenario appear in the dialog, along with their selected root directories. The root directory you selected for propagation is displayed above the **Scenarios/Hosts** table, and in the **Current Value** column marked in red.

5.  To propagate the root directory to all Replica hosts, click **OK**.

    **Note**: To exclude hosts from the root directory value propagation, clear their check boxes, and then click **OK**.

6.  After the **Propagate Value** dialog is closed, click the **Save** button on the Standard toolbar to save and apply your changes to all hosts.

# Scenario Operations

The following sections describe scenario operations:

- Save Scenarios
- Remove Scenarios
- Export Scenarios
- Import Scenarios

# Save Scenarios

There are two methods of saving scenarios, either per scenario or by a global saving of all scenarios.

**To save scenarios**

- On the Scenario pane, select the scenario and click the **Save** icon, or select the **Save** option from the **Scenario** menu.

   - or -

- Click the **Save All** icon or select **Save All** from the **Scenario** menu, to save all scenarios on the Manager.

# Remove Scenarios

**Important!** Before removing a scenario, make sure you want to permanently delete it. There is no undo action.

**To remove a scenario**

1. On the Scenario pane, select the scenario and right-click it.

   A pop-up menu appears.

2. From the pop-up menu, select the **Remove** option.

   A message appears asking you to confirm the removal.

3. Click **OK**. The scenario is permanently removed.

# Export Scenarios

You can export scenarios to other locations in order to reuse them.

**To export a scenario**

1. On the Scenario pane, select the scenario you want to export. Then, right-click it and select **Export**, or select the **Export** option From the **Scenario** menu.

   The **Exportscenario** dialog is opened.

2. Name the scenario and click the **Save** button to save it.

   The scenario is saved as a *.xmc file.

# Import Scenarios

You can import .xmc files that contain saved scenarios to your Manager. Use this option if you want to relocate scenarios from one workstation to another, or if you want to use older scenarios that were kept in your system.

**To import a scenario**

1. From the **Scenario** pane, click **Scenario Group**.

2. From the **Scenario** menu, select the **Import** option.

   An **Import scenario** dialog opens.

3. Locate the scenario you want to import and click **Open**.

   The scenario is imported to the Manager and appears on the Scenario pane.

4. Select your required options and click **OK**.

# Host Maintenance

The following sections explain the Host Maintenance option, and describe how you can prepare your hosts for maintenance procedures:

- Understanding the Host Maintenance Option
- Preparing Hosts for Maintenance Procedures

# Understanding the Host Maintenance Option

The Host Maintenance option enables you to perform maintenance procedures, such as rebooting a host or moving groups between Microsoft Cluster nodes, without performing re-synchronization once these processes are completed. Usually, when the online replication process is critically interrupted, there is a need to compare the data between the source and target hosts and make it identical, to ensure the data integrity before the replication can continue. This resynchronization process consumes time and resources. The Host Maintenance option enables you to prepare your replicated system for planned maintenance procedures and avoid resynchronization.

The hosts that can be prepared for maintenance need to participate in running scenarios. The preparation is done on one host at a time, but this host can participate in multiple scenarios. In these scenarios the host can function both as the Master and the Replica. When a host participates in a scenario that is not running, the preparation that relates to this scenario will not occur. For example, a host can participate in both File Server and Exchange scenario. If before you start preparing the host, the File Server scenario is not running, only the Exchange services will be stopped during the preparation and the Server shares will remain intact.

When the selected host functions as the Master, during the preparation process either the DB services or File shares are stopped, depending on the scenario type. Then, all changes that occurred until that moment are passed on to the Replica. Once the Replica sends to the Master an acknowledgment that all changes were applied and the data integrity is ensured, the scenario is suspended and the host is ready for maintenance. When the selected host functions as the Replica, the changes that were sent to it are applied, and the Master stops sending new changes. The new changes are saved in the meantime in the Master's spool for future update. Then, the scenario is suspended and the host is declared as ready for maintenance.

Once the maintenance procedures are completed, Arcserve RHA seamlessly resumes the real-time replication, avoiding any delay or disruption that data re-synchronization may cause.

**Important!** This option applies to Database and File Server applications. It supports both Replication and HA scenarios. However, when using this option for File Server scenarios, and you have applications that are running locally on the host you want to reboot, you need to manually stop them before starting the host maintenance preparation, and manually restart them after the maintenance is completed.

# Preparing Hosts for Maintenance Procedures

**To prepare your hosts for maintenance procedures**

1. On the Scenario pane, verify that the scenarios whose host you want to reboot are running.

   **Notes**:

   ◆ You do not have to run all the scenarios in which the host participates. The preparation will be done only on the parts that involve the running scenario, for example, Exchange services in the case of Exchange scenario.

   ◆ The host maintenance preparation cannot be performed during synchronization. If a scenario currently synchronizes, wait until it completes.

2. Click the **Launch Host Maintenance** button, or select **LaunchHost Maintenance** from the **Tools** menu.

   The **Host Maintenance** wizard opens.

   The **Host Maintenance** wizard displays all hosts that participate in the running scenarios.

   **Note**: If the same host appears under different names/IPs in different scenarios, it will appear several times in this page.

3. Select the host you want to prepare for maintenance, and click **Next**.

   The **Maintenance Procedure** page opens.

   The **Maintenance Procedure** page displays the details of the scenarios in which the selected host participates.

4. On the **Select maintenance procedure** section on the left, select the operation you want to perform and click the **Start** button.

   On the Event pane, a message appears saying: **Preparing for reboot**. Then, another message appears saying: **Ready for reboot**.

   **Note**: If a message appears saying: **Not Ready for Reboot**, it means that the preparation did not succeed, and after you reboot the host, re-synchronization will be performed.

   Simultaneously, on the Scenario pane the scenario state is changed to **Ready for HM.**

   **Note**: The scenario's state that appears on the Scenario pane refers only to the Master host's state. Therefore, if the host you are preparing for maintenance is

functioning as the Replica, you will not see its changing status in the Scenario pane, only on the Event pane and the Host Maintenance Monitor.

5. To view the status of the selected host and the scenarios in which it participates, select from the **View** menu the **Active View**, **Host Maintenance Monitor** option, or click the **Host Maintenance Monitor** button.

The **Host Maintenance Monitor** view opens.

The Host Maintenance Monitor displays all the requests for maintenance preparation. A maintenance request disappears from the Monitor when the involved scenario is either stopped or run. You cannot perform actions through the Monitor, which only displays information about the status of the current requests. The only action you can do is opening the Host Maintenance wizard by clicking anywhere in the screen and selecting **Launch Host Maintenance**.

In this Monitor, the displayed host name is its fully qualified name, and not the name under which it appears in the scenarios. All the scenarios in which this host participates appear in the Monitor.

6. After you received the message informing you that the host is ready for reboot, you can reboot your host or switch groups between cluster nodes. Once you completed your maintenance procedures, the replication process automatically resumes, without performing re-synchronization.

**Note**: If after preparing the host for maintenance, you decided not to reboot it and continue running its scenarios, you need to stop the scenarios and re-run them.

# Chapter 8: Setting Properties

This section describes how to configure scenario properties, and provides the list of the scenario properties, their corresponding values, and an explanation of each property.

# Configure Scenario Properties

Scenario Properties determine the entire scenario's default behavior concerning synchronization method, replication mode, event notification, reporting, and more.

**Notes:**

- The Properties pane and its tabs (Root Directories, Properties, Statistics) are context sensitive, and change whenever you select a different node from a scenario folder.

- Some scenario properties may be changed while the scenario is running. For more information, see the topic, Change Scenario Configuration. To change other scenario properties, you must first stop the scenario.

Each scenario is identified by its specific Product type, Server type and unique Scenario ID. The values of these items cannot be modified.

**To set or change scenario properties**

1. On the Scenario pane, select the scenario its properties you want to configure. On the Framework pane on the left, the Scenario Properties list opens.

   **Note**: A running scenario has a gray background, and scenarios that are not running have a white background.

2. If the scenario is running and the property you wish to change is one that cannot be changed while the scenario is running, click the **Stop** button on the toolbar. The scenario is stopped.

3. On the Scenario Properties list, open the desired group, select the required property, and select or enter the appropriate values. Some values can be manually entered in an edit box field, while others can be selected from a combo box or IP control by clicking the default value.

4. After you set the required properties, click the **Save** button on the Standard toolbar to save and apply your changes. Restart the scenario.

# Understanding Scenario Properties

This section lists the scenario properties[1], corresponding values, and provides an explanation for each property. The properties are listed according to their location in the respective property group:

- General
- Replication
- Event Notification
- Report Handling
- Configure Scenario Properties

---

[1]Scenario Properties are properties that affect the entire scenario. You may set general, replication, event notification, report handling and scenario properties.

# General Properties

The properties in this group cannot be changed. The Product Type and Server Type properties are set during the creation of a new scenario. The Scenario ID property is given automatically by the system. To change these properties, you need to create a new scenario.

**Product Type**

Either Replication or HA (High Availability).

**Server Type**

The type of application or database server that participates in the scenario.

**Scenario ID**

The unique ID of the scenario.

# Replication Properties

Replication includes the following properties:

▪ Mode

Arcserve RHA supports the following replication modes:
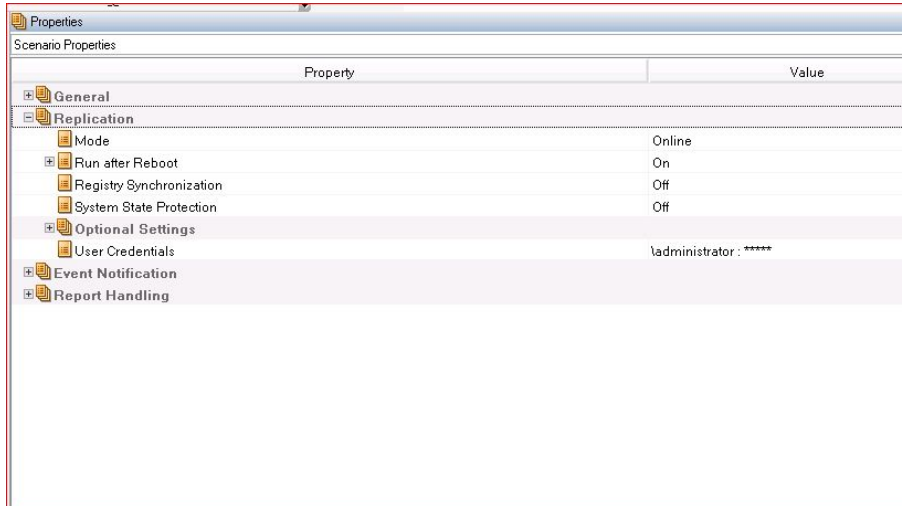


▪ Online

The Online replication mode activates real-time replication. Changes are replicated continuously, in real-time, using the XOMF driver.

The Online mode replicates all changes of files, even if files that are always open (as is the case of most database and mail servers). This mode maintains the order of file system operations. In this mode, the Engine records all I/O operations related to the root directories in journal files. The journal files are then sent to the Replicas where the operations that were recorded in the journal are replayed on the replicated files.

▪ Scheduled

Scheduled replication mode is really nothing more than an automatically performed synchronization. The synchronization can be initiated either by a manual activation, or according to a pre-defined schedule, say every few hours or once a day. This replication mode does not differ in principle from a synchronization performed as part of initializing replication. Although there is no online replication in this mode, online changes made during synchronization are replicated.

When the **Scheduling** option is selected, two options are enabled:

◆ **By User Request**

Synchronization is activated by a user running synchronization from the Manager, or PowerShell.

◆ **Weekly Activity**

In the Weekly Activity setting of the Scheduled mode, servers are synchronized at a scheduled, fixed time. When you select this option, you need to set the schedule for the recurring synchronization.

For a detailed description of synchronization scheduling, refer to Schedule Synchronization.

◆ Periodic Replication

File changes are not replicated continuously, but periodically aggregated. Aggregated changes are then propagated to the Replica on a scheduled basis. In Scenario Properties, under Replication, expand Mode and set the Schedule and Schedule Setting properties.

You may experience delays when applying replication data depending on the data size and number of files to apply, because the process assures data consistency. You can set the parameter, KeepUndoForPeriodic, to false to halt generation of the undo file and speed up the application of aggregated changes, but should not do so for database scenarios. This parameter is in the ws_rep.cfg file.

▪ Run after Reboot

If the Master is rebooted, Arcserve RHA automatically re-synchronizes the Master and the Replica after the reboot.

◆ Automatic Synchronization

Synchronization ensures that a set of folders and files on a Replica server involved in a scenario is identical to the set on the Master.

◆ Synchronization Type

– **FileSynchronization**

Compares files on the Master and Replica servers, and whenever they are different, copies the entire missing or modified files from Master to Replica.

In order to synchronize your data for the first time, you need to choose the File Synchronization mode for each scenario. In subsequent cases, this synchronization type is best suited for File Server (a large number of small and medium sized files),

preferably with the checked **Ignore Files of Same Size/Time** option. This may significantly decrease synchronization time.

— **Block Synchronization**

Performs a block-by-block comparison of the Master and Replica files, and copies over only those blocks that are different. When differences exist between files, instead of requiring the transfer of the entire file, the block synchronization transfers only changes.

Block synchronization is the method appropriate for database applications, such as MS Exchange, Oracle or SQL Server. You should use this method while clearing the **Ignore Files of Same Size/Time** option (unless database files are closed on the Master server).

◆ Ignore Files of Same Size/Time

Skips comparison of files with the same path, name, size and modification time. This assumes that the files are identical.

Select this option when you are absolutely sure that files of this type are indeed identical. This option is best suited for File Server scenarios. It is not appropriate for database files for such applications as Exchange, SQL, or Oracle, since these databases modify files (leaving them open) without changing file modification time. You can use this option in database scenarios only when synchronized databases are un-mounted and files are closed on the Master server.

This option can reduce the overall synchronization time dramatically, but you must remember that this comes at the expense of content verification.

▪ Registry Synchronization

When set to On, this option enables you to synchronize the Master and the Replica registry keys on a scheduled basis.

For a detailed description of registry synchronization, refer to Synchronize Registry Keys.

▪ System State Protection

When set to On, this option enables you to save snapshots of the system and boot files of the Master on the Replica. To activate this option, you need to set the System State Protection schedule, and to define which Replica host will store the snapshots. For a detailed description, refer to Protecting Your System State.

- Optional Settings

    - Replicate NTFS Compress Attribute

      (For Windows only) Replicates compress attribute of files or directories during synchronization and replication.

    - Replicate ACL

      Replicates ACLs for files and directories during synchronization and replication.

      Retain Local Account Names (On) **(For Windows only)**

      RHA engine retains the local name in the Replica machine. Before you run the scenario, manually create the same local users or groups in both Master and Replica. Use this property for workgroup environments.

    - Synchronize Windows Shares

      If a directory has been set to allow sharing, then setting this Share option to On duplicates the property in the replicated directory. This occurs only during synchronization and on Windows OS.

    - Keep the Archive Attribute on Replica

      If the master and replica files are identical, then do not change the archive attribute on replica during synchronization.

    - Prevent Automatic Re-sync upon Error

      An uncommon critical error on the Master can stop the replication continuance. In this case, setting this option to On prevents automatic re-synchronization. When this option is Off, re-synchronization starts automatically upon an error occurrence.

    - Stop the scenario when a disk/spool is full

      For scenarios configured with multiple Replica servers, this property lets you choose to stop only the affected Replica or the entire scenario when a spool or disk on any Replica is full. Set this property to Off to stop only the affected Replica and On to stop the entire scenario.

    - Number of Streams

      This property lets you more effectively use bandwidth in a high latency (WAN) environment. You can use the GUI or the ws_rep.cfg file to set this property to a value between 1 and 10. The default value uses a single stream, but is over-ridden by any value set in the ws_rep.cfg file. Any other value set in the GUI overrides the ws_rep.cfg setting.

The property opens a number of streams to send and receive data equal to the value you specify, and can be used with Bandwidth Scheduling.

**Important**! In LAN environments, you should use one stream.

- User Credentials

Lets you enter the user credentials to access the host folder. If the user has no permission to access the root directory, then scenario will not run.

# Event Notification Properties

**Notification**

When an event occurs, you can set the system to run a script, send an email notification, or write it to Windows event log.

**Notify by Email**

Defines whether to send the details of an event by email to a specified address. If several events occur immediately one after the other, the system aggregates them and sends their details in one email.

- **On Error Only** - Set to On to receive email notifications when the application detects errors.

- **Mail Server**

  Enter the mail server hostname or IP.

- **Mail Server Port**

  Enter the port number for the mail server.

- **Authentication Settings**

  Click to open the Mail Authentication Settings dialog and enter mail authentication details such as user name, password, and proxy settings.

- **Email Address - To**

  Enter the receiver email address.

- **Email Address - From**

  Enter the sender email address.

**Execute Script**

Specifies a script to run whenever an event occurs.

- **Script Name (full path)**

  Enter the name and full path of the script that is invoked once an event occurs.

- **Arguments**

  Additional arguments to pass to the script, which is specified in the previous property. Any arguments entered here follow the argument sent automatically by Arcserve RHA, which include the event details written in a notification file. Arguments entered here are static values.

**Note**: On Windows x64, you cannot run scripts that activate applications with UI.

**Write to Event Log**

Writes the events to the Windows event log.

# Report Handling Properties

**Report Saving**

Enter the report saving settings.

**Report Directory**

Specifies the location where the reports are saved.

**Report Retention (days)**

Specifies the number of days to retain replication reports. The default is Unlimited.

**Notify by Email**

Defines whether to send reports by email to the specified address.

- ✦ **Mail Server**

  Enter the mail server hostname or IP.

- ✦ **Mail Server Port**

  Enter the port number for the mail server.

- ✦ **Authentication Settings**

  Click to open the Mail Authentication Settings dialog and enter mail authentication details such as user name, password, and proxy settings.

- ✦ **Email Address - To**

  Enter the receiver email address.

- ✦ **Email Address - From**

  Enter the sender email address.

**Execute Script**

Specify a script for Arcserve RHA to run whenever it generates a report.

- ✦ **Script Name (full path)**

  Enter the name and full path of the script that is invoked once a report is generated.

- ✦ **Arguments**

  Additional arguments to pass to the script specified in the previous property. Any arguments entered here follow the argument sent automatically by Arcserve RHA. This argument defines the full path of the generated report file and its type. Arguments entered here are static values.

**Note**: On Windows x64, you cannot run scripts that activate applications with UI.

# Schedule Synchronization

When selecting a scheduled replication mode it means that synchronization will be performed automatically on a regular basis. Once you select this option the following flexible scheduling capabilities are offered:

- Synchronization on selected days of the week and for specific hours in a 24-hour cycle.

- Synchronization over selected periods (for example, once every 36 hours) in a 7-day cycle.

- Exclusion of specific dates.

**To open the schedule**

1. On the Scenario Properties list, open the **Replication** group. On the **Mode** property, select the **Scheduling** value.

   The **Schedule** option appears.

2. On the **Schedule** option, select the **Weekly Activity** value. Then, on the **Schedule Settings** property, click the **Not Set** value.

   The **Schedule Setting** dialog appears.

3. Set the schedule for automatic synchronization according to the guidelines described in the following sections.

   The following figure shows a sample scheduling setting in which the solid blue rectangles indicate the days and hours scheduled for synchronization. The excluded dates area lists the specific dates on which synchronization is not performed.

This section contains the following topics:

- Set a Schedule for Automatic Synchronization
- Exclude Dates from Scheduled Synchronization
- Set Advanced Schedules

# Set a Schedule for Automatic Synchronization

The following instructions describe how to set and clear hours and days in the **Schedule Setting** dialog for automatic synchronization.

**To set a specific hour/day**

- Select a single rectangle for a specific hour/day. Click the **Set** button to actually mark and enable that hour/day.

**To set a specific hour for each day in the week**

- Select a column and click the **Set** button.

**To set each hour in a specific day in the week**

- Select a row and click the **Set** button.

**To set a repetitive cycle**

- Enter a valid number of hours in the **Every** box, and click the **Apply** button.

**Note**: More than one rectangle can be simultaneously set by clicking and dragging the mouse. You can also use the **Ctrl** and **Shift** keys to set several dates at once.

**To clear a setting**

- Use the same technique of selecting, and click the **Clear** button.

**Important!** If synchronization is running, and the next scheduled synchronization hour comes up, the new synchronization stops the current one and starts again from the beginning.

# Exclude Dates from Scheduled Synchronization

You can set specific dates that will be excluded from the automatic synchronization.

**To exclude dates from automatic synchronization**

- In the **Schedule Setting** dialog, on the **Excluded dates** section select or enter the excluded date in the **dates** box. Then, click the **Add** button.

The selected date appears in the **Excluded dates** list.

**To remove an excluded date**

- On the **Excluded dates** list, select the entry, and click the **Remove** button. Multiple entries can also be selected by dragging the mouse over them.

# Set Advanced Schedules

The **Advanced Schedule Setting** page allows you to set times that are not on the hour.

**To open the Advanced Schedule Setting page**

- In the **Schedule Setting** dialog, click the **Advanced** button at the bottom.

**To return to the Standard Schedule Setting dialog**

- In the **Advanced Schedule Setting** page, click the **Standard** button at the bottom.

# Setting Master and Replica Properties

This section describes how to configure the Master and Replica properties, and provides the list of their properties, the corresponding values, and an explanation of each property.

**Related Topics**

- Configure Master or Replica Server Properties
- Understand Master Properties
- Understand Replica Properties
- Schedule the Bandwidth Limit
- Propagate Property Values

# Configure Master or Replica Server Properties

To configure Master or Replica properties, the scenario must be stopped.

**To set Master or Replica properties**

1. On the Scenario pane, select the Master or Replica whose properties you want to configure.

   On the Framework pane on the right, the Master/Replica Properties list opens.

   **Note**: A running scenario has a gray background, and scenarios that are not running have a white background.

2. If the scenario is running, click the **Stop** button on the toolbar. The scenario is stopped.

3. On the Master/Replica Properties list, open the desired group, select the required property, and select or enter the appropriate values. Some values can be manually entered in an edit box field, while other values can be selected from a combo box or IP control by clicking the default value.

4. Click the **Save** button on the toolbar to save and apply your changes.

# Understanding Master Properties

This section lists the [Master properties](#)[1], corresponding values, and provides an explanation for each property.

**Note**: On Windows x64 systems, you cannot run scripts that activate applications with a graphical user interface.

The properties are listed according to their location in the respective property group:

- Host Connection
- Replication
- Spool
- Event Notification
- Reports

---

[1]Master Properties are properties that are set on and control the behavior of the Master server in a scenario. You may set general, replication, event notification, report handling and scenario properties.

# Host Connection Properties on the Master

**Management IP Address**

Enter the IP address of the Master host. If the Master name is changed, the IP address is updated. The Master can also be changed by entering another IP address in this field.

**Replication IP address**

Enter the replication IP address of the master host. If the replication IP is provided, the engine uses it to transfer data during synchronization and replication instead of the management IP address.

**Port Number**

Enter the number of the incoming port that are used for TCP communications. This port number can be changed to any unused port. Since the Engine uses only one port, make sure that the Engine uses the port that is specified here. The default port number is 25000.

**HTTP Proxy Server Configuration**

Allows data replication using an HTTP proxy server. To use a proxy server, configure the HTTP proxy server on the master.

# Replication Properties on the Master

**Run Script before Synchronization**

Triggers a script to run before each synchronization. The synchronization process does not start until this script run is completed.

**Script Name**

Enter the full name and path of the script.

**Arguments**

Arguments to pass to the script that is specified in the previous property. Arguments are static values.

**Run Script after Synchronization**

Triggers a script to run after each synchronization. The synchronization process does not wait for this script run to finish.

**Script Name**

Enter the full name and path of the script.

**Arguments**

Arguments to pass to the script that specified in the previous property. Arguments are static values.

**Compress Data during Transfer**

Compresses data before sending it to the Replica. This option optimizes bandwidth and transfer time. If the Master host is a busy production server, we recommend activating this feature on the first-level Replica that resides on the same LAN, and not on the Master.

◆ Compress data is resource consuming, and impacts server performance. If the typical file format being transferred does not compress much, this option is a waste of processing power and time. Although the transmission bandwidth can be lessened, the overall replication time is a function of compressibility and available power.

◆ Already compressed files such as .zip, .rar, .gz, .jpeg, etc., and any small file whose size is less than 512 bytes, are not compressed.

**IO Throttling During Synchronization**

Enables you to control I/O throttling during synchronization.

**Enable Scheduled Bookmarks (Only for DB Scenarios)**

This property lets you create periodic bookmarks for DB scenarios according to the schedule you set. Use the generated bookmarks in Rewind for periodic recover points. This feature in only for DB scenarios such as MSSQL and Exchange scenarios.

**Default:** Every two hours

### Run Script upon Trigger File Creation

[For File Server only] Defines whether special actions should be triggered via a script, when a specified trigger file appears.

#### Trigger File Name

The name of the file that triggers the script, which is specified in the next property. The script is triggered once the file creation event occurs.

#### Script to Run

### Script Name

This script is invoked, when the trigger file that is specified in the previous property is created. Enter the full name and path of the script.

### Arguments

Arguments to be passed to the script specified in the previous property. Arguments must be static values.

### User Credentials

Lets you enter the administrator credential, or the proper credentials in the ACL group if ACL license is applied to access the host folder. If the authentication fails, then the scenario does not run.

# Spool Properties

The spool is a folder on a disk where data to be replicated is temporarily stored (that is, spooled). The spool stores changes captured during replication for some period of time before applying them to the Replica server. The spool parameters determine 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.

**Max Spool Size**

Enter the maximum spool size allowed. This disk space is used only if needed - it is not pre-allocated. The default is Unlimited. To enter a value of **Unlimited**, enter a zero.

**Min Disk Free Size**

Enter the free disk space threshold at which the system issues an error and stops replication.

**Spool Directory**

Enter the directory to be used to store the spool. The default directory is *INSTALLDIR/tmp* on Windows.

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

# Event Notification Properties

**Notification**

When an event occurs, you can set the system to run a script, send an email notification, or write it to the Windows event log.

**Notify by Email**

Defines whether to send the details of an event by email to a specified address. If several events occur immediately one after the other, the system aggregates them and sends their details in one email.

- ✦ **On Error Only -** Set this option to On to receive emails when the application detects errors.

- ✦ **Mail Server**

  Enter the mail server hostname or IP.

- ✦ **Mail Server Port**

  Enter the port number for the mail server.

- ✦ **Authentication Settings**

  Click to open the Mail Authentication Settings dialog and enter mail authentication details such as user name, password, and proxy settings.

- ✦ **Email Address - To**

  Enter the receiver email address.

- ✦ **Email Address - From**

  Enter the sender email address.

**Execute Script**

Specifies a script for Arcserve RHA to run whenever it sends a report.

- ✦ **Script Name (full path)**

  Enter the name and full path of the script that is invoked once an event occurs.

- ✦ **Arguments**

  Additional arguments to pass to the script, which is specified in the previous property. Any arguments entered here follow the argument sent automatically by Arcserve RHA, which include the event details written in a notification file. Arguments entered here are static values.

**Write to Event Log**

Writes the events to the Windows event log.

# Reports Properties

**Generate Synchronization Report**

Specifies whether to generate a synchronization report.

**Generate Detailed Report**

Specifies whether to generate a detailed synchronization report.

**Generate Replication Report**

Specifies whether to generate a replication report. Since replication is continuous, specify the frequency of the report generation in the property below.

**Generation Frequency (hours)**

Specifies how often to generate the replication report.

**Generate Detailed Report**

Specifies whether to generate a detailed replication report.

**Report Handling**

**Notify by Email**

Specify whether to send reports by email to the specified address.

◆ **Mail Server**

Enter the mail server hostname or IP.

◆ **Mail Server Port**

Enter the port number for the mail server.

◆ **Authentication Settings**

Click to open the Mail Authentication Settings dialog and enter mail authentication details such as user name, password, and proxy settings.

◆ **Email Address - To**

Enter the receiver email address.

◆ **Email Address - From**

Enter the sender email address.

**Execute Script**

Specify a script for Arcserve RHA to run whenever it sends a report.

◆ **Script Name (full path)**

Enter the name and full path of the script that is invoked once a report is generated.

◆ **Arguments**

Additional arguments to pass to the script specified in the previous property. Any arguments entered here follow the argument sent automatically by Arcserve RHA. This argument defines the full path of the generated report file and its type. Arguments entered here are static values.

**Write to Event Log**

Writes the events to the Windows event log.

# Understanding Replica Properties

This section lists the Replica properties[1], corresponding values, and provides an explanation of each property.

**Note**: On Windows x64 systems, you cannot run scripts that activate applications with a graphical user interface.

The properties are listed according to their location in the respective property group:

- Host Connection
- Replication
- Virtual Machine
- Spool
- Cloud
- Recovery
- Volume Snapshot Management Properties
- Scheduled Tasks
- Event Notification
- Reports

**Related Topics:**

- Configure Master or Replica Properties

---

[1]Replica Properties are properties that are set on and control the behavior of the Replica server in a scenario. You may set host connection, replication, spool, recovery, scheduled tasks, event notification and report handling properties.

# Host Connection Properties on the Replica

**Management IP Address**

Enter the IP address of the Replica host. If the host name is changed, the IP address is updated. The host can also be changed by entering another IP address in this field.

**Replication IP Address**

Enter the replication IP address of the Replica host. If the replication IP is provided, the engine uses it to transfer data during synchronization and replication instead of the management IP address.

**Port Number**

Enter the number of the incoming port used for TCP communications. It can be changed to any unused port. Since the Engine can use only one port, make sure that the Engine uses the port specified here. The default port number is 25000.

**Enable HTTP Proxy Server**

Allows data replication using an HTTP proxy server. To use a proxy server, configure the HTTP proxy server on the master.

# Replication Properties on the Replica

**Run Script before Synchronization**

Triggers a script to run before each synchronization. The synchronization process does not start until this script run is completed.

**Script Name**

Enter the full name and path of the script.

**Arguments**

Arguments to pass to the script specified in the previous property. Arguments are static values.

**Run Script after Synchronization**

Triggers a script to run after each synchronization. The synchronization process does not wait for the script run to finish.

**Script Name**

Enter the full name and path of the script.

**Arguments**

Arguments to pass to the script specified in the previous property. Arguments are static values.

**Compress Data during Transfer**

Compresses data before sending it to the Replica. This option optimizes bandwidth and transfer time.

◆ Compress data is resource consuming, and impacts server performance. If the typical file format being transferred does not compress much, this option is a waste of processing power and time. Although the transmission bandwidth can be lessened, the overall replication time is a function of compressibility and available power.

◆ Already compressed files such as .zip, .rar, .gz, .jpeg, etc., and any small file whose size is less than 512 bytes, are not compressed.

**Encrypt Data During Transfer**

Encrypts data so that data sent between the Replica and its parent node are transferred securely (also called Secure Communication). This property can only be set on the Replica host, which means this property will not exist on the Master. However, during recovery or when you run backward scenarios, this

option will be adjusted so that the link between the two hosts remains secure (or non-secure, for plain data, if this option is set to No Encryption).

**Note:** See Change the Engine Security Method to define your own security parameters to encrypt the data.

**Keep Deleted Files during Synchronization**

During synchronization, do not remove from the Replica files that were deleted from the Master. Best suited for cases in which several scenarios use the same Replica directories.

**Keep Deleted Files during Replication**

During replication, do not remove from the Replica files that were deleted from the Master.

**Bandwidth Limit (Kbps)**

Controls the size of the allowed incoming bandwidth on the Replica host. You can either define one limit size that will apply to all hours of the day, or you can specify different values for different hours. The default value is **Unlimited.**

For a detailed description of bandwidth scheduling, refer to Schedule the Bandwidth Limit.

**Stop Database on Run**

When set to On, if a database scenario (Exchange, SQL, Oracle) is running and the database is running on the Replica server, Arcserve RHA stops the database services before running the scenario. [Does not apply to HA scenarios]

**Store System State on this Replica**

This option can be enabled only when the **System State Protection** property in the Scenario Properties list is set to On. For more information, refer to Protecting Your System State.

**Retry if File is Busy**

These options are relevant only for Windows servers. If changes were received for a busy file (opened as non-shared for read), these options define how many times and at what interval to attempt replacing this file with the one that contains the changes.

**Number of Attempts**

Enter the number of attempts to be made for replacing a modified file that is busy (and therefore cannot be replicated). If the file is not released before the last attempt is made, the change is lost and an error message is initiated.

**Interval between Attempts (msec)**

The time between an unsuccessful attempt and the next attempt.

**Run Script upon Trigger File Creation**

[For File Server only] Defines whether special actions should be triggered via a script, when a specified trigger file appears.

**Trigger File Name**

Enter the name of the file that triggers the script, which is specified in the next property. The script is triggered once the file creation event occurs.

**Script to Run**

◆ **Script Name**

This script is invoked, when the trigger file specified in the previous property is created. Enter the full name and path of the script.

◆ **Arguments**

Arguments to be passed to the script specified in the previous property. Arguments must be static values.

**User Credentials**

Lets you enter the user credentials that has full control permission to parent directory available in the root directory. If you do not have such permission, then the scenario does not run.

# Change the Engine Security Method

Arcserve RHA Engine uses a predefined security method. However, if you set the Replica property Encrypt Data During Transfer to *On*, you can also change the default SSL self-signed certificate, RSA private key and cipher list to use your own security parameters in the Engine configuration file. The Engine configuration file that contains the security method is **ws_rep.cfg**.

**To change the Engine security method**

1. If there are running scenarios that are currently using the Engines for which you want to change the security method, stop them through Arcserve RHA Manager.

2. Log in to the Master and Replica host where the Engine is running.

3. In the Services dialog, stop the Engine service on both the Master and Replica servers.

4. Using Windows Explorer, browse to Engine installation directory, where the ws_rep.cfg file is located.

   **Note**: The default installation directory is: *C:\Program Files\CA\Arcserve RHA\Engine*.

5. Open the ws_rep.cfg file with WordPad or another text editor.

   **Note**: We do not recommend using Notepad, due to its limited view options.

6. Do the following in the ws_rep.cfg file:

   a. Find the # SSLSelfSignedCertificate = "[INSTALLDIR]/cacert.pem" section.

   b. Change the SSLSelfSignedCertificate = "[INSTALLDIR]/cacert.pem" to reflect the name of the SSL self-signed certificate that you want to use and remove the # symbol at the beginning of the line.

   c. Find the # SSLRSAPrivateKey = "[INSTALLDIR]/cakey.pem" section.

   d. Change the SSLRSAPrivateKey = "[INSTALLDIR]/cakey.pem" to reflect the name of RSA private key that you want to use and remove the # symbol at the beginning of the line.

7. Save the ws_rep.cfg file.

   **Important!** While the configuration files on the Master and Replica servers can be different, you must ensure that the parameters you use to change the security method be identical in both the Master and Replica servers' ws_rep.cfg file.

The Engine security method is changed in the ws_rep.cfg file.

8.  Start the Engine Service on both the Master and Replica.

9.  Open the Manager, highlight the scenario, and restart it.

**Note:** If the SSL self-signed certificate and RSA private key fail to load, the default setting is used and a warning message will display in the Arcserve RHA Manager.

# Virtual Machine Properties

When you are working with Cloud scenarios, in addition to <u>Cloud Properties</u>, you can also review and manage Virtual Machine properties:

**Virtual Platform Setting**

Lets you review the settings for the following read-only properties:

**Virtual Platform Type**

Identifies the virtual platform type of the cloud account.

**Virtual Platform**

Identifies the virtual platform server of the cloud account.

**Port**

Identifies the port number used to connect to the virtual machine.

**SSL Connection**

Identifies whether the SSL (secure socket layer) connection is on or off.

**Virtual Machine Setting**

Lets you define the following properties:

**EC2 Instance Type**

Lets you specify the size assigned to the EC2 instance on the virtual machine. You can specify the appropriate instance type based on the operating system of the master and the requirements of your environment. Instance type options include:

- Small Instance
- Large Instance
- Extra Large Instance
- High-Memory Extra Large Instance
- High-Memory Double Extra Large Instance
- High-Memory Quadruple Extra Large Instance
- High-CPU Medium Instance
- High-CPU Extra Large Instance

Available options are related to the master's platform. If the master is a 32-bit operating system, only the Small Instance and High-CPU Medium Instance

are available. If the master is a 64-bit operating system, then all of the other types are available.

**Virtual Machine Name**

Specifies the name of the virtual machine managed on the virtual platform server.

# Spool Properties

The spool parameters determine 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.

**Max Spool Size**

Enter the maximum spool size allowed. This disk space is used only if needed - it is not pre-allocated. The default is Unlimited. To enter a value of **Unlimited**, enter a zero.

**Min Disk Free Size**

Enter the free disk space threshold at which the system issues an error and stops replication.

**Spool Directory**

Enter the directory to be used to store the spool. The default directory is *INSTALLDIR/tmp* on Windows.

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

# Cloud Properties

Cloud includes the following properties:

**Cloud Provider**

Identifies the name of the cloud service running the selected cloud instance. This is a read-only property.

**Cloud Account ID**

Identifies the account ID of the AWS account. This is a read-only property.

**Cloud Region**

Identifies the VPC region of the AWS account. This is a read-only property.

**Cloud Instance ID**

Identifies the ID of the cloud instance. This is a read-only property.

**Cleanup Cloud Resources When Remove Scenario**

Lets you specify whether to clean up cloud resources when a scenario is removed. For Full System EC2 Data Replication or High Availability scenarios, several cloud resources can be used such as the cloud instance used for fail-over, volumes, and snapshots. If these cloud resources are useless after a scenario is removed, you can enable this option to clean up these resources. This option is disabled by default.

**Shutdown Instance on scenario stop**

Lets you specify whether to shut down the replica instance automatically on scenario stop. This option is disabled by default, which means that the replica instance will not be automatically stopped if the scenario is stopped.

# How to Stop Scenario When Spool is Full

If you have scenarios configured with multiple Replica servers, the property, Stop scenario if any replica spool/disk is full, lets you choose between stopping only the affected Replica or the entire scenario when a spool or disk is full. The default value is On, which instructs the application to stop the entire scenario when the spool or disk on any replica is full. When this property is set to Off, only the Replica is stopped. No changes are sent to the Replica until it is resumed. When the Replica resumes, resynchronization for only that affected Replica is triggered.

Set the property from the Replication, Optional Settings group on the Scenario Properties tab.
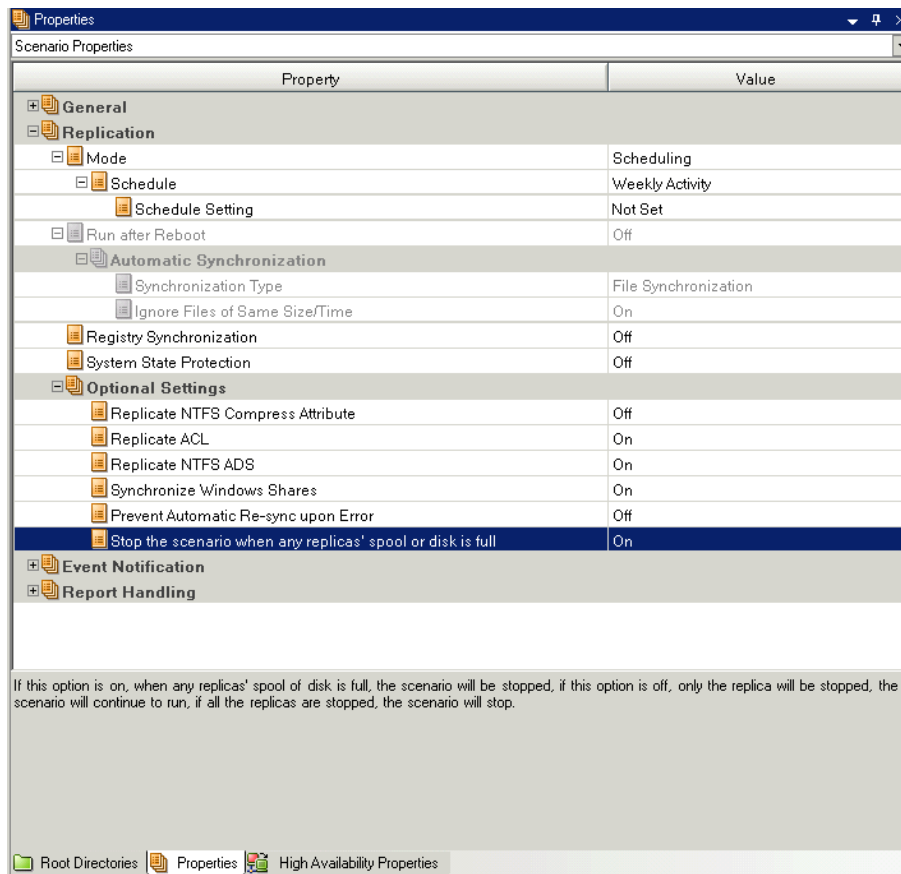


Arcserve RHA logs warnings that the spool limit has been exceeded, or the disk is almost out of space in the Events pane of the Manager. Depending on how you set the property, the scenario or the Replica is then stopped, allowing you to clear disk space. The spool is automatically cleared.

To start a stopped Replica, right-click it from the Manager and choose Start the Replica from the shortcut menu. Resynchronization is triggered and replication resumes after resynchronization completes.

# Recovery Properties

**Replication Delay**

Data replication can be delayed in the Replica server spool before sending it to this Replica. This is useful against data corruption or viruses. It enables stopping replication before corrupted or infected data is written to the Replica.

### Delay Interval (min)

Enter the number of minutes of the replication delay.

**Data Rewind**

Keeps undo information needed to recover data from a certain action or point in time. It is useful in cases in which corrupted data on the Master was replicated to the Replica, and you want to restore the data to its previous state before the corruption occurred. Data Rewind is enabled for online replication only.

### Retention Period (min)

I/O operations are recorded in the Rewind journal for this number of minutes. Then, they are discarded in first-in-first-out order.

### Max Disk Size (MB)

Enter the maximum disk space allocated for the Rewind journal. Once this size is reached, old records are discarded in FIFO order.

# Volume Snapshot Management

Enter the number of snapshots you want to keep, storage volume, and maximum storage size.

# Scheduled Tasks Properties

**Suspend**

Refer to Schedule Replication Suspension.

**Replica Integrity Testing for Assured Recovery**

Refer to Setting Assured Recovery Properties.

# Event Notification Properties

**Notification**

When an event occurs, you can set the system to run a script, send an email notification, or write it to Windows event log.

**Notify by Email**

Defines whether to send the details of an event by email to a specified address. If several events occur immediately one after the other, the system aggregates them and sends their details in one email.

- ◆ **On Error Only**

    Set this property to On to receive emails when the application detects errors.

- ◆ **Mail Server**

    Enter the mail server hostname or IP.

- ◆ **Mail Server Port**

    Enter the port number for the mail server.

- ◆ **Authentication Settings**

    Click to open the Mail Authentication Settings dialog and enter mail authentication details such as user name, password, and proxy settings.

- ◆ **Email Address - To**

    Enter the receiver email address.

- ◆ **Email Address - From**

    Enter the sender email address.

**Execute Script**

Specifies a script for Arcserve RHA to run whenever it sends a report.

- ◆ **Script Name (full path)**

    Enter the name and full path of the script that is invoked once an event occurs.

- ◆ **Arguments**

    Additional arguments to pass to the script, which is specified in the previous property. Any arguments entered here follow the argument sent automatically by Arcserve RHA, which include the event details written in a notification file. Arguments entered here are static values.

**Write to Event Log**

Writes the events to the Windows event log.

# Reports Properties

**Generate Replication Report**

Specifies whether to generate a replication report. Since replication is continuous, specify the frequency of the report generation in the property below.

**Generation Frequency (hours)**

Specifies how often to generate the replication report.

**Generate Detailed Report**

Specifies whether to generate a detailed replicationreport.

**Generate Assured Recovery Report**

Specifies whether to generate Assured Recovery report.

**Report Handling**

**Notify by Email**

Specify whether to send reports by email to the specified address.

- **Mail Server**

  Enter the mail server hostname or IP.

- **Mail Server Port**

  Enter the port number for the mail server.

- **Authentication Settings**

  Click to open the Mail Authentication Settings dialog and enter mail authentication details such as user name, password, and proxy settings.

- **Email Address - To**

  Enter the receiver email address.

- **Email Address - From**

  Enter the sender email address.

**Execute Script**

Specify a script for Arcserve RHA to run whenever it sends a report.

- **Script Name (full path)**

  Enter the name and full path of the script that is invoked once a report is generated.

◆ **Arguments**

Additional arguments to pass to the script specified in the previous property. Any arguments entered here follow the argument sent automatically by Arcserve RHA. This argument defines the full path of the generated report file and its type. Arguments entered here are static values.

# Schedule the Bandwidth Limit

Arcserve RHA enables you to control the size of the allowed incoming bandwidth on the Replica host. For each day of the week, you can either define one limit size that will apply to all hours of the day, or you can specify different values for different hours. By using the Bandwidth Scheduler, you can decrease the bandwidth size on busy hours and increase it on off-peak hours in order to optimize your bandwidth resources.

You can also perform offline synchronization. For more information, see the topic, Synchronization Methods.

**Notes**:

- The bandwidth limit that you set for one Replica host does not apply to other Replica hosts that reside in the same replication tree. You need to change each Replica host definition separately.

- The default value for the Bandwidth Limit option is "**Unlimited**". It means that no restriction is imposed on the bandwidth between the Master and the Replica.

**To schedule the bandwidth limit:**

1. On the Replica Properties list, open the **Replication** group. On the **Bandwidth Limit** property, click the Value box that contains the **Unlimited** default value.

   The **Bandwidth Scheduler** dialog opens.

2. Set the daily schedule for incoming bandwidth size according to the following guidelines:

   - On the **Day/Hours** table, select one or several rectangles for the period in the day for which you want to define a certain bandwidth size.

     **Note**: You can set more than one rectangle simultaneously by clicking and dragging the mouse. You can also use the **Ctrl** and **Shift** keys to set several dates at once.

   - Once the rectangles are marked, from the **Bandwidth Values** section click the values (in Kbps) you want to apply on the selected hours.

     The rectangles of the selected hours now contain the value you selected.

3. To apply similar bandwidth to all days, apply the bandwidth to Sunday and click **Apply the Sunday's value to all work days** orrepeat the above process

for all hours. If no size is defined for a specific hour, the **Unlimited** default value is applied to it.

   **Note:** To clear the setting, click the **Reset** button.

4.  Once you finished defining the Bandwidth Scheduler, click **OK** to save your setting and close the dialog.

   The setting you defined now appears in the **Bandwidth Limit** value box in the Properties list.

5.  To save your setting, click the **Save** button on the Standard toolbar.

# Propagating Property Values

Arcserve RHA enables you to apply the values you set for one scenario to multiple scenarios at once. Instead of separately configuring the properties of each scenario, you can propagate the value of one scenario to as many scenarios as you want. A good example would be to use this option to simultaneously change the e-mail notification address for multiple scenarios. You can propagate the values of scenarios, Master hosts and Replica hosts.

**Note**: The scenarios must be stopped to apply property changes.

**To propagate values of properties**

1. On the Scenario pane, select either the scenario, the Master or the Replica whose properties you want to propagate. On the Framework pane on the left, the Properties list opens.

2. On the Properties list, open the desired group and right-click the property value you want to propagate. The **Propagate Value** pop up command opens.
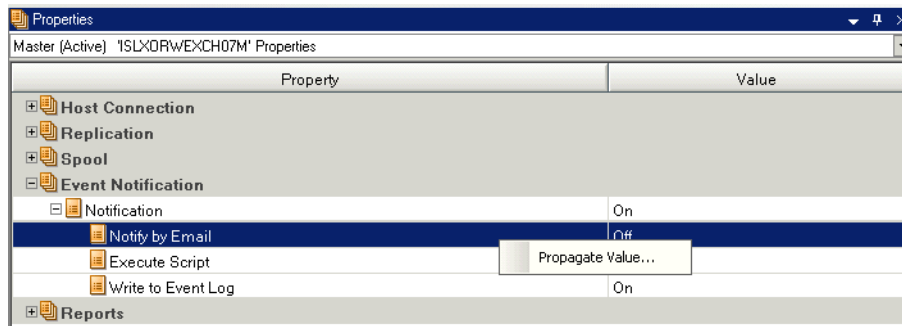


3. Click the **Propagate Value** command. The **Propagate Value** dialog opens.

   All scenarios in your Manager appear in the dialog, while the scenario whose property value you want to propagate is marked in red. The property and the value you can propagate are displayed above the **Scenarios** table, and in the **Current Value** column.

4. To propagate the value to all scenarios, click **OK**.

   **Note**: To exclude scenarios or hosts from the value propagation, clear their check boxes, and then click **OK**.

5. After the **Propagate Value** dialog is closed, click the **Save All**  button on the Standard toolbar to save and apply your changes to all scenarios.

# Change Configuration when a Scenario is Running

You can change certain properties while the scenario is running, allowing you to troubleshoot scenario problems without having to first halt replication.

- Spool properties
    - Max Pool Size
    - Min Disk Free Size
    - Spool Directory
- Script properties
    - Script Name
    - Arguments
- Is alive properties
    - Is Alive Timeout
    - Heartbeat Frequency
- Bandwidth limit
- I/O Throttling During Synchronization
- Schedule properties
    - Periodic Replication Frequency
    - Generate Replication Report and Frequency
    - Mode, Schedule, Schedule Setting
    - Registry Synchronization and Synchronization Frequency
- Event Notification properties
    - Notify by Email: On error only, Mail Server, Email Address From, To
    - Execute Script
    - Write to Event Log

**Change these properties using the Arcserve RHA Manager.**

1. From the Scenario list, click the server to configure.

2. Click the Properties tab and browse for the property you wish to change.

3. Set the new value, as desired.

4. Click Apply, Discard changes, or Cancel as required. Click View change details to expand the dialog to show the original and new values of the properties you modified.

During editing, the green arrow (indicating a running scenario) changes to a pencil icon to indicate editing. Arcserve RHA distributes the changed scenario files to the participating hosts and verifies the changes. If a host cannot receive the file or verification fails, the changes are not applied.

# Protect Your System State

System State Protection allows you to include critical system-related components in your replication and high availability scenarios, so that you can recover these components after a failure. System State Protection is server-type independent, which means you can enable this feature in any Arcserve RHA scenario. When System State Protection is enabled, information about the snapshot schedule and file management is stored directly in the scenario configuration file. Snapshot files are then transferred to all Replica servers configured under the Master, including any "child" Replicas, if its Retain System Snapshot property (Transferring Path) is enabled.

**Note**: Even if a child Replica's Retain System Snapshot property is enabled, a child Replica will not receive system state snapshots if this property on its parent is disabled.

System State Protection lets you:

- Configure System State Protection in the Scenario Creation Wizard
- Set periodic backup intervals
- Send the snapshot file to more than one replica
- Restore from a System State snapshot

The following components are always included when System State Protection is enabled:

- boot files
- COM+ Class Registration database
- Registry service

The following components are also included in System State Protection depending upon the operating system:

- Windows Server 2003 Operating Systems

    – All files protected by WFP

    – performance counter configurations

    – Active Directory (ADSI) on systems that are domain controllers

    – SYSVOL directory replicated by File Replication Service (FRS) on systems that are domain controllers

– Certificate server on systems that provide Certificate Authority

– Cluster database on systems that are a node of a Windows cluster

▪ Windows Server 2008 Operating Systems

– Certificate Services database

– Active Directory Domain Services (NTDS)

– SYSVOL directory (FRS Writer)

– Cluster Service information

– Microsoft Internet Information Services (IIS) meta-directory (IIS Meta-base Writer/IIS Configuration Writer)

– System files under WFP (System Writer)

**Note**: For Windows Server 2003 and higher, system state snapshots are taken using System Writer. Refer to the Microsoft website for more information on Backing Up and Restoring System State Under VSS.

**Related Topics:**

▪ How to Configure System State Protection

▪ Modify Scenario System State Protection

▪ Restore System State Data

▪ Command Line Enhancements for System State Protection

▪ Additional System State Information

# How to Configure System State Protection

By default, System State Protection is set to Off. You can enable System State Protection during scenario creation using the wizard, or you can enable it for existing scenarios using the Scenario Properties pane of the Arcserve RHA Manager.

For either method, you need to perform the following tasks:

- Enable the System State Protection property for the scenario

- Set the snapshot schedule, as desired.

- Enable the Store System State Protection property on one or all Replicas involved in the scenario.

**Note**: After System State is restored on a Replica, that machine must be rebooted so the system state can take effect. You can set the Reboot After Recovery property to On, if desired.

# Configure System State Protection in the Scenario Creation Wizard

System State Protection can be enabled for any Arcserve RHA supported server type, from directly inside the Scenario Creation Wizard.

**To enable System State Protection in the Wizard**

1. From the Arcserve RHA Manager, start the scenario creation wizard using the toolbar button or the menu command: New, Scenario.

2. Complete the steps in the usual manner for the server type you selected until you reach the Scenario Properties dialog.

3. Under Replication properties, set System State Protection to On.

4. Under System State Protection, click the Value for Set Schedule to access the Schedule Setting dialog.

5. Set the schedule for taking system state snapshots.

6. Complete scenario creation steps as usual until you reach the Master and Replica Properties dialog.

7. Expand Replication properties on the Replica and set the Store System State on this Replica property to On. You may also set additional storage properties at this time. For more information, refer to the topic, Store System State Properties.

8. Save and run the scenario. System State Protection starts.

# Configure System State Protection for Existing Scenarios

If you did not enable System State Protection when creating the Scenario, you can configure the System State Protection property outside the Scenario Creation Wizard using the following procedure.

**Note:** By default System State Protection is set to off.

Before performing this procedure, stop the scenario. From the <cawan> Manager, select the scenario and click the Stop button from the toolbar.

**To configure system state protection for existing scenarios**

1. From the <cawan> Manager select a Scenario to change its System State Protection property.

2. On the Scenario Properties panel, expand the Replication node and set the Enable System State Protection node property to On.

If the Replica is not configured, a message dialog opens.

1. Click OK.

2. Expand the System State Protection node and set a schedule, if desired. For more information, refer to the topic, Set the System State Protection Schedule.

3. On the Replica, enable the store system state protection property. For more information, refer to the topic, Configure System State Protection on the Replica.

4. Save the Scenario.

# Set the System State Protection Schedule

If you did not enable set a System State Protection schedule when creating the Scenario, you can configure the System State Protection schedule property outside the Scenario Creation Wizard using the following procedure.

**Note:** By default System State Protection schedule is set to off.

**To set the System State Protection Schedule**

1. From the Arcserve RHA> Manager, select a Scenario to set the System State Protection Schedule property.

2. On the Scenario Properties panel, expand the Replication node and expand the System State Protection node.

3. Click the Value column of the Set Schedule node.

    The Set Schedule dialog opens.

4. Set the day, start time, frequency, and exclude dates as needed.

5. Click OK.

# Configure System State Protection on the Replica

If you did not enable System State Protection on the Replica when creating the Scenario, you can configure the Store System State on this replica property outside the Scenario Creation Wizard using the following procedure.

**Notes:**

- You can enable System State Protection on multiple Replicas.

- By default System State Protection is set to off.

**To configure system state protection of the Replica**

1.  From the Arcserve RHA Manager, select a Replica to enable its System State Protection property.

2.  On the Properties panel, expand the Replication node and set the Store System State on this replica node property to On.

3.  Modify the Replica storage property values as required. For more information, refer to the topic, Store System State Protection Properties.

4.  Save the Scenario.

# Store System State Protection Properties

You may set the following properties on the Replica server to manage system state snapshot storage:

**Copies to keep**

Specify the number of System State snapshots to keep on the Replica. The default value is eight. For an unlimited number of snapshots, enter a zero. If the number of snapshots exceeds the set value, the oldest snapshots are deleted to make space for subsequent snapshots.

**Max Total Disk Size (MB)**

Specify the total amount of disk space to allocate for System State snapshots.

Default Values by Operating System:

- Windows 2003: 8192

- Windows 2008: 16,384

For an unlimited amount of space, enter a zero. If snapshots fill the space allocated, the oldest snapshots are deleted to make space for subsequent snapshots.

**Note**: Arcserve RHA checks the Copies to Keep and Max Total Disk Size settings periodically, not at scenario start.

**Min Disk Free Size (MB)**

Specify the minimum amount of free disk space for System State snapshots. The default value is 1024. If the free disk space is less than the set value, the oldest snapshots are deleted to ensure that the minimum free disk space value is maintained.

**Directory**

Specify the directory where the System State snapshot is stored.

# Modify Scenario System State Protection

When a scenario is stopped, you may modify its properties, including System State Protection. The changes you make take effect when the scenario is restarted.

**System State Protection Property**

If you disable the System State Protection property on an existing scenario, you are prompted to delete existing snapshots. If you select Yes, all snapshots from the Replica are deleted when the scenario is restarted. If you select No, all snapshots are retained.

**Store System State on this Replica Property**

If you disable the Store System State on this Replica property on an existing scenario, you are prompted to delete existing snapshots. Choose Yes to delete all snapshots or choose No to retain them.

You can modify all properties under the Store System State on this Replica group:

**Copies to Keep**

You can adjust the number of snapshots to keep.

**Max Total Disk Size (MB)**

You can adjust the maximum total disk size setting.

**Min Disk Free Size (MB)**

You can adjust the minimum disk free size to which snapshots are saved. If you set this value to 0, size is unlimited. If this setting is less than the size when the scenario is running, the oldest snapshot is deleted to make space for a new one.

**Modify Directory to Store Snapshots**

You can change the directory to which stored snapshots are saved. If you specify an invalid path, you are notified that snapshots cannot be saved. If a valid path is set, you are prompted to move old snapshots to a new location. If you select No, old snapshots are deleted.

For more information, refer to the topic, Store System State Protection Properties.

# Restore System State Data

The recovery process for restoring System State data is similar to the usual recovery process, with one distinction. if System State Protection is enabled, you are prompted at recovery time to select a Recovery Source (see the following procedure).

Arcserve RHA waits for application data recovery to complete first, including synchronization, before recovering the system state. You must select a Replica to act as the restore source machine. During the restore process, a new scenario is created in which the Master and Replica servers reverse roles. The scenario then transfers the System State snapshot to the original Master.

**To restore system state data**

1. From the Arcserve RHA Manager select a Scenario to restore its System State data.

2. Select the Replica host.

3. From the toolbar, click the Restore Data  button.

   The Recovery Source dialog opens if the scenario has System State Protection enabled.

4. Click Restore Application Data and Restore System State, and then click Next.

   **Note**: If you select Restore System State only, you will not be able to choose an application recovery point. If you disable System State Protection, you will not see the Recovery Source screen.

   The Rewind Point Selection dialog opens.

5. From the Recovery Point Selection dialog, click Select Rewind Point to open the Recovery Point Selection dialog.

6. Set the following criteria:

   **Time**

   Select any previous backup point from the list.

   **Recovery to**

   Recovery data to the default location or browse to a specific location.

7. Click Finish.

8. Reboot the Replica now.

# Command Line Enhancements for System State Protection

The following commands have been added to PowerShell to support System State Protection:

**set-properties scenario_name index value**

Use the set-properties command to set System State Protection for a scenario.

To obtain index values, use the get-properties command.

**set-hostproperty scenario_name replica_name index value**

Use the set-hostproperty command to enable the Store System State property on a Replica.

To obtain index values, use the get-hostproperties command.

**RecoveryMode [A|S|B]**

Use A to recover application data only. (Default setting)

Use S to recover system state only.

Use B to recover both.

**RebootAfterRecovery [0|1]**

Use 0 to skip reboot (default)

Use 1 to enable Master reboot after recovery

# Additional System State Information

**System State Protection in High Availability Scenarios**

> After switchover, system state snapshots are not retained on the original master.

# Chapter 9: Recovering Data and Servers

This section describes how to recover a server, restore lost data using the Manager, how to set bookmarks, and how to rewind data.

This section contains the following topics:

- [The Data Recovery Process](#)
- [Recover Lost Data from Replica](#)
- [Setting Bookmarks](#)
- [Data Rewind](#)

# 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.

# Recover Lost Data from Replica

You can restore data from a replica server. Doing so may require login credentials for the machine you select.

**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 opens.

   **Notes**:

   - 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.

   - The **Include Registry Keys synchronization** checkbox is enabled, only if you activated the Registry Synchronization property before starting the scenario. If the checkbox is enabled, you can select it to include the synchronized Registry Keys in the recovery process.

5. Click **Next**. The **Synchronization Method** page opens.

6. Select the appropriate Synchronization Method, based on scenario type. 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.

   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 on 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 Bookmarkoption 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.**
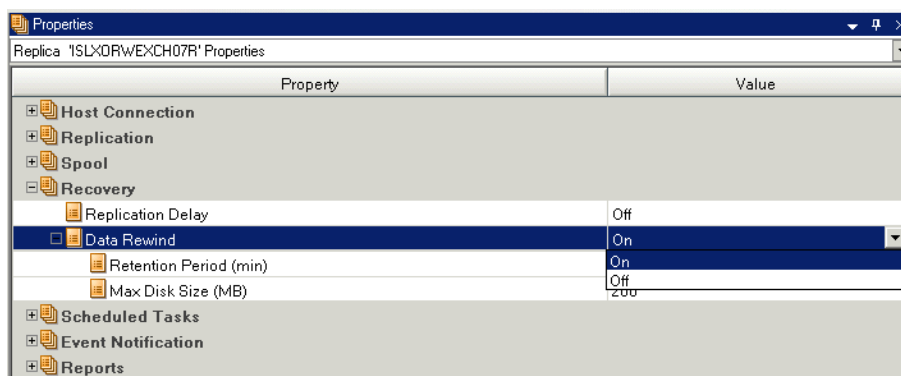


If this option is set to Off, the system will not register data rewind points.

**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. You cannot rewind to another point past the time when you made changes to Replica files.

**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. In the following example, a File Server scenario is used, but the steps are the same for all scenario types.

**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. If you are prompted for user credentials, enter the appropriate inform-
ation and click OK.

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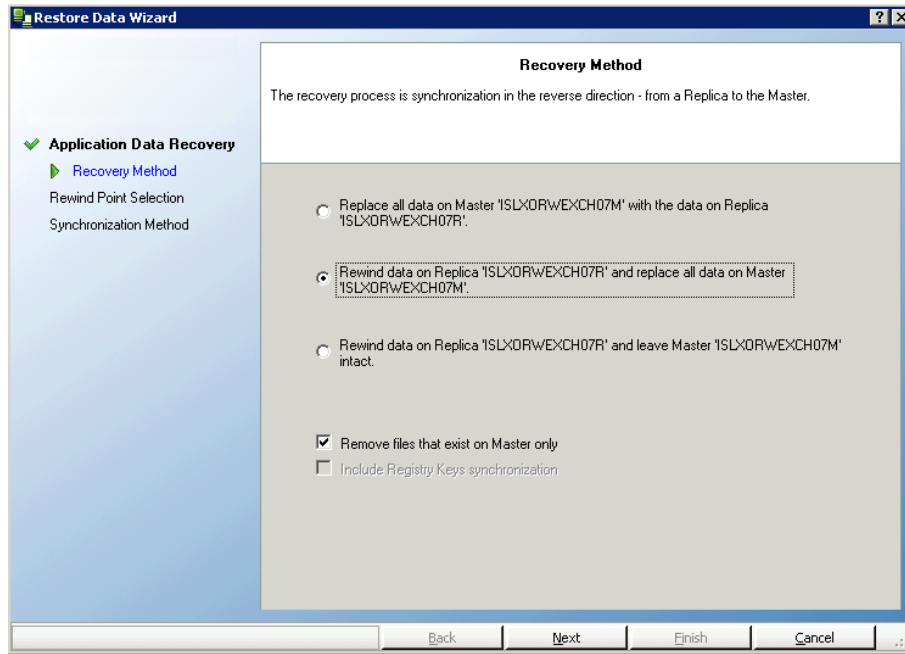


5. Select one of the Rewind data options, depending on whether you want the
rewind data synchronized back to the Master (option 2) or left on the Replica
only (option 3).

**Notes**:

◆ 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.

◆ The **Include Registry Keys synchronization** checkbox is enabled, only if
you activated the Registry Synchronization property before starting the
scenario. If the checkbox is enabled, you can select it to include the syn-
chronized Registry Keys in the recovery process.

After you select a Rewind data option, a Recovery scenario is automatically cre-
ated. 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 files that were automatically registered by the system and user-defined book-marks.

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

   **Notes:**

   - If the **Select Rewind Points** dialog is empty, make sure that the Data Rewind property is enabled.

   - The entire list can be exported to an Excel file by clicking the **Export to Excel** button on the bottom-left corner.

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 appropriate synhronization 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 cre-dentials** 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 gen-erated.

Now, the Replication process can restart on the original scenario.

# Restore Full Systems

As an Administrator, you are responsible for ensuring that data is replicated and failover happens in case the Master server fails. On a failover or switchover, you can recover the entire data to an active server from the Replica.

The following diagram illustrates how the recovery process restores data to an active server.
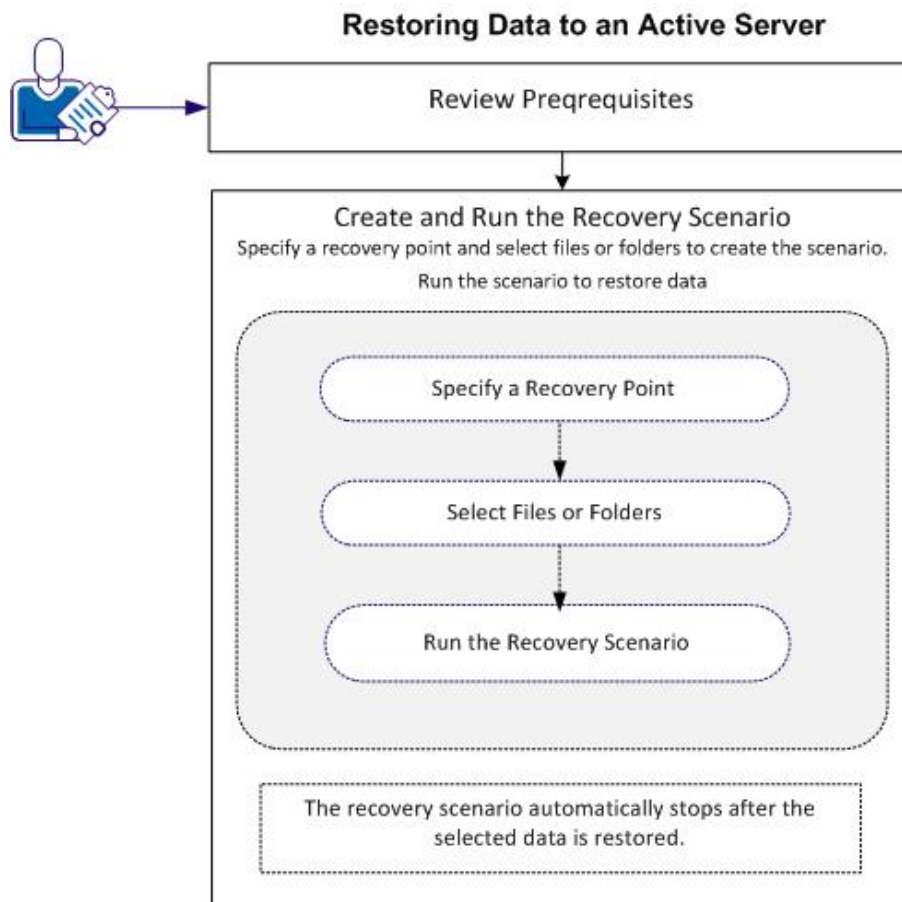


Perform the following tasks to restore data:

- Review Prerequisites
- Create and Run the Recovery Scenario

# Review Prerequisites

Before you can restore data to an active server, consider the following points:

- Start recovery only when the full system scenario is stopped or a switchover-/failover is triggered.

- Restore data to a similar machine. Install the same version of the operating system and applications as much as possible.

- Build the new Master first and then launch the recovery process.

- For an FSP scenario, do not manually stop the virtual machine that performed a switchover or failover before you start the Master.

- Stop critical services before launching the recovery process, such as Database services and Exchange server services. When the services are running, some applications lock their files and cannot be opened during the recovery process. Start these services again after the recovery.

# Create and Run the Recovery Scenario

Review the prerequisites carefully before you launch the recovery. Use the Data Recovery wizard to restore data from the Replica.

**Follow these steps:**

1. Log in to Arcserve RHA as an Administrator.

2. From the Quick Start pane, click Scenario Management to open the Arcserve RHA Manager.

3. Launch the Data Recovery Wizard from the Manager screen.

4. Enter the credentials, when prompted.

5. Select a recovery point on the Recovery Point Selection screen. For example, click Latest System Status and then click Next.

   The virtual machine shuts down.

6. On the Master Root Directories screen, expand the source host and then, click or clear check boxes to include or exclude folders from the list. If you wish, you can apply a file filter. Click Next.

   Important: Do not select the C:\Windows folder as it can lead to certain conflicts later, especially if the operating system versions are different. If you do select the folder, then system registry files such as Sam, Security, Software, and Default are not restored.

   **Note:** We recommend skipping the RHA Engine folder (C:\Program Files\CA\Arcserve RHA\Engine) when you are not sure about the engine versions.
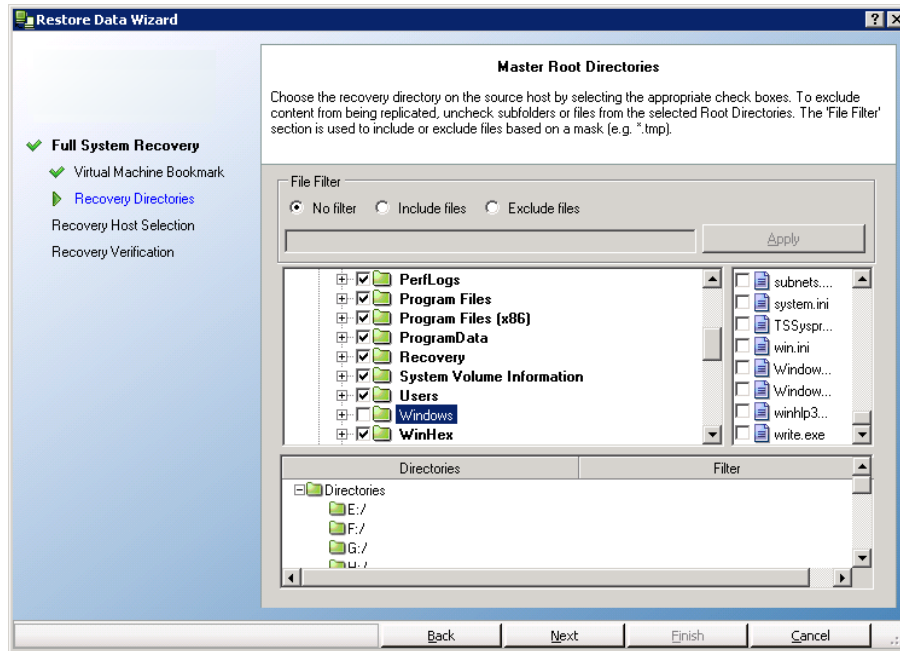
7. On the Recovery Host screen, enter the required credentials and click Next.

   Wait while the verification completes.
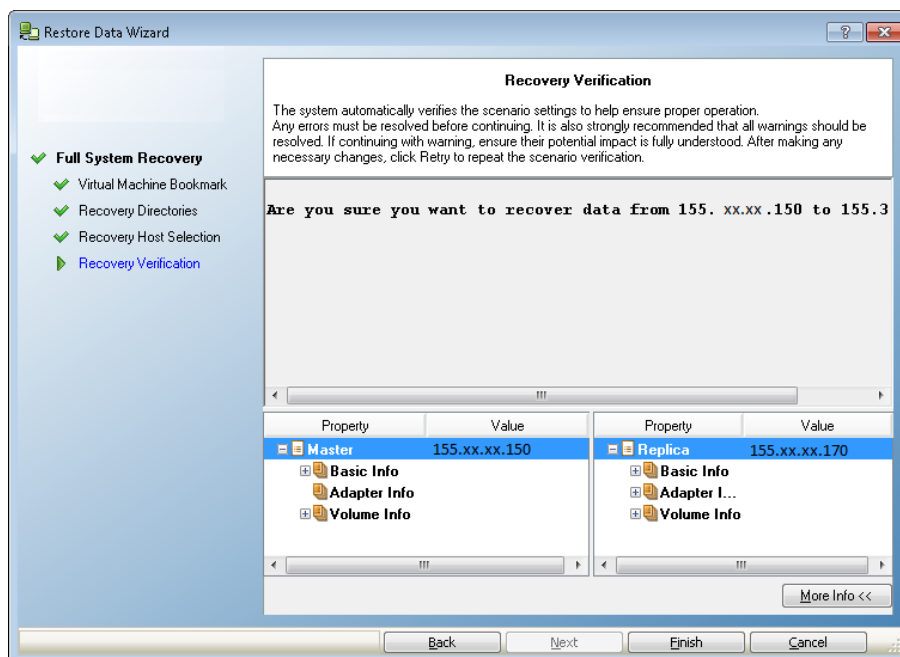


8. When the Recovery Verification screen opens, review the information or click More Info to see more details. Click Finish.

   A new scenario, Recovery_<Scenario Name>, is added to the list on the Manager screen. View the statistics for progress. The new scenario automatically stops once the selected data is recovered. The system reboots if necessary.

# Chapter 10: Switching Over and Switching Back

This section explains the High Availability process and the switchover and switch-back procedures, and describes the following operations: creating HA scenario, performing switchover and switchback, and recovering the active server.

This section contains the following topics:

- [Switchover](#)
- [Switchback](#)
- [Recovering Active Server](#)
- [Understanding the High Availability System and the Switchover and Switch-back Procedures](#)
- [Setting High Availability Properties](#)

# Switchover

This section contains the following topics:

- The Switchover Workflow
- Initiate Switchover

# The Switchover Workflow

Switchover (or failover) is the process of changing roles between the Master and Replica, meaning making the Master server the standby server, and the Replica server the active server.

Switchover can be triggered automatically by Arcserve RHA when it detects that the Master is unavailable (failover). Alternatively, Arcserve RHA can simply alert you to the problem, and then you can manually initiate switchover from the Manager.

During the creation of HA scenario, you define how you want the switchover to be initiated. If you selected in the **Switchover and Reverse Replication Initiation** page the **Initiate Switchover manually** option, you need to perform a manual switchover. However, 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, for example, 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.

If you choose to initiate switchover automatically, after the Master is considered to be down, Arcserve RHA automatically tries to restore the services and databases on it to their active state. First, Arcserve RHA tries to restart services that it previously checked to be managed. If the services are running, it then tries to mount the databases. If all attempts fail, Arcserve RHA initiates failover. These attempts to restore the services and databases are not performed if the switchover is initiated manually.

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

# Initiate Switchover

**To initiate a 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. Optionally, select **Do no wait until all journals are applied** to immediately perform switchover even before all journals are applied. If you do not select this check box, the switchover process gets initiated only after all journals are applied.

4. Click Yes.

5. Depending on your scenario configuration, the **Run a Reverse Replication Scenario after the Switchover** check box is either selected or cleared. You can change your default configuration only for the switchover you are about to perform, by selecting or clearing the check box. Then, click **Yes** on the **Perform a 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.

6. After the switchover is completed the scenario stops.

   **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 original Master becomes the Replica and the original Replica becomes the Master.

**Important**! If the Master server crashes or is rebooted during a switchover, the process stops. If this happens, you may need to recover the active server.

# Switchback

This section contains the following topics:

- The Switchback Workflow
- Initiate Switchback

# The Switchback Workflow

After a switchover was initiated, whether manually or automatically, at some point you will want to reverse back the server roles, and make the original Master the active server again and the Replica the standby server. Before you switch back the roles between them, if you want the data on the active server, meaning the original Replica, to overwrite the data on the standby server, you need to run a reverse scenario (also called "backward scenario").

During the creation of the HA scenario, you defined how you want the reverse scenario to be initiated. If you selected the **Initiate Reverse Replication automatically** option, replication in the reverse direction (from Replica to Master) automatically begins after a switchover, once the original Master will become available. However, if you selected the **Initiate Reverse Replication manually** option, you need to perform a manual switchback. If the manual option is selected and you will not initiate a manual switchback, a resynchronization of the data from Replica to Master will have to be performed, even after a test of clean switchover without an actual Master failure.

**Note**: After a switchover, in certain circumstances you may want to switch the Master and Replica roles without overwriting the original Master data with the Replica data. To perform this, use the **Recover Active Server** option.

# Initiate Switchback

**To initiate a switchback**

1. Ensure that both Master and Replica servers are available on the network and that the Engine is running.

2. Open the Manager and select the desired scenario from the Scenario pane.

3. [Skip this step if the backward scenario is already running, and move to step 7.]

   If the backward scenario is not running, select the **Run** button to start the scenario. Arcserve RHA detects that a switchover has occurred, verify its state and configuration, and prompts you to approve the running of the backward scenario.

   **Note**: The **Advanced** button opens an additional pane with detailed information about the hosts that participate in the scenario.

4. Click the **Run** button to start the backward scenario.

   The **Run** dialog opens.

5. For Microsoft Exchange, select **Block Synchronization.** For File Server, click File Synchronization and click **OK**. The resynchronization starts.

   Wait until the resynchronization is completed.

6. Once the resynchronization is completed, you receive the following message in the Event pane: **All modifications during synchronization period are replicated**. Then, replication from the active server to the standby server starts.

7. Now, you can reverse back the roles between the Master and Replica servers. To reverse the roles, while the backward scenario is running, click the **Perform Switchover** button, or select the **Perform Switchover** option from the **Tools** menu.

   A confirmation message opens.

8. Click **Yes** on the **Perform Switchover** confirmation dialog. This procedure initiates a switchback from the original Replica server to the Master server.

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

Now, you can run again the scenario in its original (forward) state.

# Recovering Active Server

In certain circumstances, it may be necessary to intentionally make the Master or Replica server the active server without completing the synchronization process. This may happen when, for example, a switchover occurred but no data was changed on the Replica server, and you might even have newer data on the Master server. In this case, it is undesirable to synchronize data from the Replica to the Master server. Arcserve RHA allows you to manually select the active server option through a process called **Recover Active Server**.

There might be cases when the switchover process did not complete correctly, but using the **Recover Active Server** option does not resolve the problem, or that you prefer not to use this option in order to correct the situation. In such cases, you can manually recover the active server from outside the Manager. The type of procedure you will need to perform depends on the type of redirection method you used for the switchover.

The available options for recover active server following incomplete switchover are as follows:

- Use the Recover Active Server option from the Manager.
- Manually recover the active server from outside the Manager.

# Recover Active Server Using the Manager

If the switchover process did not complete properly, Arcserve RHA enables you to manually select which server will act as the active server through a process called **Recover Active Server**.

**Important!** While this option is the right choice in certain 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, which then overwrites what may be a more current data set. When using the **Recover Active Server** option, Arcserve RHA redirects 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.

**To recover active server using the Manager**

1. On the Scenario pane, select the scenario whose active server you want to recover and stop it.

2. From the **Tools** menu, select the **Recover Active Server** option.

   Arcserve RHA verifies which host is currently the active one, and presents the results in the **Recover Active Server** dialog.

3. Click either the **Make Master Active** or **Make Replica Active** button, depending on which server you want to assume the active role.

   Messages appear in the Event pane, informing you that one of the server becomes active while the other becomes inactive. Now, the host you selected becomes the active server, and users are directed to it.

**Important!** In a disaster situation, if a standard switchover occurs and users are redirected to the Replica server for any period of time, it is important to replicate all changes that occurred on the Replica back to the Master, before making the Master server active again. Using the **Recover Active Server** option in such a situation may result in loss of data.

# Recover Active Server from Outside the Manager

If the switchover process does not complete correctly for some reason, and using the **Recover Active Server** option from the Manager does not resolve the issue, you can try one of the following manual tasks appropriate to the redirection method you used:

- If **Move IP** redirection method is used, manually remove the additional IP from the Master.

- If **Switch Computer Name** redirection method is used, manually switch the Computer Names between the Master and the Replica.

- If both **Move IP** and **Switch Computer Name** redirection methods are used, manually remove the additional IP from the Master, and switch the Computer Names between the Master and Replica.

# Manually Recover a Failed Server - Move IP Address

**To recover a failed server when Move IP redirection is used**

1. Boot the Master server without a network connection, to avoid IP conflicts.

   The Is Alive check from the Replica to the Master recognizes resource failure, and makes the original Replica the active server.

2. From the **Advanced TCP/IP Settings** dialog on the Master, remove the additional IP address.

3. Reboot the Master server, and reconnect to the network.

4. If the HA scenario is not already running, start the scenario from the Manager by clicking the **Run** button.

   If the **Run Reverse Replication Scenario after Switchover** property was set to On, the scenario runs in backward mode, from the original Replica to the original Master. The original Master which now acts as the standby server.

5. Wait for the synchronization to complete.

6. To return the active role to the Master server, perform a manual switchover by clicking the **Perform Switchover** button on the Standard toolbar.

   **Note:** We recommend that you perform the manual switchover outside of normal business hours.

# Manually Recover a Failed Server - Switch Computer Name

**To manually recover a failed server when the Switch Computer Name redirection method is used**

1. Boot the Master server without a network connection, to avoid duplicate network names.

   The Is Alive check from the Replica to the Master recognizes resource failure, and makes the original Replica the active server.

2. Rename the Master server to <NewServerName>-RHA, and move it to a temporary workgroup.

   For example, if the server is called "Server1", rename it to "Server1-RHA".

3. You are required to reboot this machine.

   After the reboot completes, the following error appears: **At least one Service could not be started**. Ignore this message, which it is normal under these circumstances, because the Engine usually runs in a domain account.

4. Connect to the network.

5. Rejoin the domain, ensuring that you use the -RHA name assigned in step 2.

6. Reboot the machine.

7. If the HA scenario is not already running, start the scenario from the Manager by clicking the **Run** button on the Standard toolbar.

   If you selected the **Run Reverse Replication Scenario after Switchover** option, the scenario runs in backward mode, from the original Replica to the original Master. The original Master now acts as the standby server.

8. Wait for the synchronization to complete.

9. To return the active role to the Master server, perform a manual switchover by clicking the **Perform Switchover** button on the Standard toolbar.

   **Note:** We recommend that you perform the manual switchover outside of normal business hours.

# Manually Recover a Failed Server - Move IP and Switch Computer Name

**To manually recover a failed server when both IP and Switch Computer Name Redirection are used**

1. Repair any hardware problems that could have caused the switchover issue, if any.

2. Reboot the server without a network connection to prevent IP conflicts.

   The Is Alive check from the Replica to the Master recognizes resource failure, and makes the original Replica the active server.

3. From the **Advanced TCP/IP Settings** dialog on the Master, remove the additional IP address.

4. From the **System Properties** dialog, **Computer Name** tab, 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. After the reboots completes, reconnect to the network. An error message appears: **At least one service failed during system startup.** Ignore this message, which it is normal under these circumstances, because the Engine usually runs in a domain account.

7. Rejoin the domain, making sure you use the -RHA name, and reboot again.

8. If the HA scenario is not already running, start the scenario from the Manager by clicking the **Run** button on the Standard toolbar. If you selected the **Run Reverse Replication Scenario after Switchover** option, the scenario runs in backward mode, from the original Replica to the original Master. The original Master now acts as the standby server.

9. Wait until the synchronization completes.

10. To return the active role to the Master server, perform a manual switchover by clicking the **Perform Switchover** button on the Standard toolbar.

   **Note:** We recommend that you perform the manual switchover outside of normal business hours.

# Understanding the High Availability System and the Switchover and Switchback Procedures

A High Availability scenario incorporates all the functionality and workflow of a replication scenario, but it adds three important new elements: pre-run verification, monitoring of the Master and the application running on it, and the switchover process itself.

■ **Pre-run verification**

During a switchover, there are many things that can go wrong - there might be problems with permissions, or with the application configuration, or even with the settings within the HA scenario itself. For this reason, when HA scenario is created and initiated, Arcserve RHA performs an extensive list of checks. These checks are designed to determine, whether any of the common issues that are known to cause problems during switchover can be found. When such issues are found in the pre-run verification, errors and warnings are presented, prompting you to solve these issues before running the HA scenario.

■ **Automatic monitoring of the Master and the application running on it**

As soon as the scenario is initiated, the Replica checks the Master on a regular basis, by default every 30 seconds. There are three types of monitoring checks - a ping request that is sent to the Master in order to verify that the Master is accessible and alive; a database check that verifies that the appropriate services are running and the data is in good state; a user-defined check that can be tailored 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. Depending on the HA scenario configuration, this will cause Arcserve RHA to send you an alert or to automatically initiate a switchover.

■ **Switchover and switchback workflow**

In an initial HA scenario, the Master is the active computer, and the Replica is the standby computer. The standby computer is continuously checking the state of the active one, to determine whether it is alive and to decide whether to assume the active role.

A switchover can be triggered automatically or with the push of a button. The first time a switchover occurs, the Replica that was on standby becomes the active computer, and the Master reverts to a standby mode (assuming it is still

operational). When the Master (now the 'standby') is ready, a switchback process can be initiated, either automatically or manually. Following the switchback, the Master again becomes active, and the Replica returns to its previous standby and monitoring role.

**Note:** After a connection loss, during the attempt to reconnect, a node (either Master or Replica) tries to determine its role. If the two nodes establish themselves as Masters, upon reconnection the newest active Master will continue to act as the Master, while the older one will turn into the standby Replica.

**Important**! After switchover, the "Server" service on the standby server, used to support file, print, named-pipe sharing, becomes inaccessible for ten minutes after switchover occurs. See the option, HASharesAccessTimeout, in the ws_rep.cfg file.

# Setting High Availability Properties

This section describes how to configure High Availability properties, and provides a list of the HA properties, their corresponding values, and an explanation of each property.

**Note**: These options are available only to those who are licensed for High Availability.

Some properties apply only to certain server types ( Exchange, SQL, etc.). For more information, see the specific Operation Guide.

**Related Topics:**

- [Configure High Availability Properties](#)
- [Active and Standby Hosts](#)
- [Move IP Redirection](#)

# Configure High Availability Properties

The HA property values determine the entire scenario's default behavior concerning network traffic redirection method, database management, and more.

**Notes:**

▪ The Properties pane and its tabs (Root Directories, Properties, Statistics) are context sensitive. The displayed content is changed whenever you select a different node from the scenario folder.

▪ To configure scenario properties, the scenario must be stopped.

**To set scenario properties**

1. On the Scenario pane, select the HA scenario whose properties you want to configure.

2. On the Framework pane on the left, select the High Availability Properties tab.

   The High Availability Properties list opens.

   **Note**: A running scenario has a gray background, and scenarios that are not running have a white background.

3. If the scenario is running, click the **Stop** button on the toolbar. The scenario is stopped.

4. On the Scenario Properties list, open the desired group, select the required property, and select or enter the appropriate values. Some values can be manually entered in an edit box field, while others can be selected from a combo box by clicking the default value.

5. After you set the required properties, click the **Save** button on the Standard toolbar to save and apply your changes.

# Understanding High Availability Properties

This section lists the High Availability properties, corresponding values, and provides an explanation of each property. It also explains the Active and Standby host concepts. The HA properties are listed according to their location in the respective property group:

- Switchover

- Hosts

- Network Traffic Redirection

- Is Alive

- DB Management/Application/Shares Management

- Actions upon Success

# Switchover

The system continuously checks for a switchover situation, as defined in the Is Alive properties, and informs the user according to the defined notification settings. When the system detects a switchover situation, the **Perform Switchover** Automatically option is checked to determine if there should be an automatic switchover, or only a notification of the situation. In the latter case, switchover may be triggered with the **Perform Switchover** menu command or toolbar button.

**Switchover Hostname**

Enter the hostname/IP address of the Replica host to which the Master switches. You can choose only one Replica.

Any time this Name/IP, or Master's Name/IP, is changed, all the switchover properties are reset to their default values.

**Perform Switchover Automatically**

When this option is On, switchover is initiated automatically if the Master server is down or database failure detected.

**Run Reverse Replication Scenario after Switchover**

After a switchover, this option determines whether replication in the reverse direction would begin automatically.

When set to On, Arcserve RHA verifies that the data integrity is kept. If the data is found to be consistent, re-synchronization is avoided and the backward scenario is initiated. However, in DB scenarios when this option is set to On, the reverse replication starts in a suspended state. It is unsuspended only after the database on the Active server has passed all tests typically performed in the **Check DB** properties in the **Is Alive** properties.

# Hosts

**Master Fully Qualified Name**

Indicates the fully qualified name of the Master host. It is obtained automatically from the operating system and you cannot change it manually.

**Replica Fully Qualified Name**

Indicates the fully qualified name of the Replica host. It is obtained automatically from the operating system and you cannot change it manually.

# Network Traffic Redirection

There are three redirection methods:

- Move IP
- Redirect DNS
- Switch Computer Name
- Switch Computer Alias

You may also specify User-defined scripts for each server type.

The three redirection methods can be used according to the following table:

|  | Move IP | Redirect DNS | Switch Computer Name | Switch Alias Name |
|---|---|---|---|---|
| Microsoft Dynamics | No | Yes | Yes | Yes |
| File Server | Yes | Yes | Yes | Yes |
| Full System HA | No | Yes | No | No |
| Exchange | Yes | Yes | No | Yes |
| Oracle | Yes | Yes | No | Yes |
| SQL | Yes | Yes | Yes | Yes |
| IIS | Yes | Yes | Yes | Yes |
| Control Service | No | Yes | Yes | Yes |
| SharePoint | Yes | Yes | Yes | Yes |
| VMware vCenter | Yes | Yes | Yes | Yes |
| Hyper-V | No | Yes | No | No |

**Note:** For computer alias names, Arcserve RHA supports UNC access.

After a switchover takes place, the A-records of Master and Replica servers are changed on the DNS server, depending on which Network Redirection Methods you select. The following table shows what network redirection methods impact the DNS A-record.

- If the Network Redirection method is MoveIP, the DNS A-record is not impacted.
- If the Network Redirection is DNS redirection, the A-record Master redirects to the Replica IP after switchover, while the Replica A-record doesn't change.
- If the Network Redirection is Switch Computer Name, the Master A-record changes to Master-RHA after switchover, and the Replica A-record changes to Master.

- If the network redirections are DNS redirection and switch computer name, the result is same as only DNS redirection.

If you want to keep the Replica A-Record, use the DNS Redirection method or DNS and Switch Computer Name methods.

**Move IP**

During switchover, the switchover IP(s) are released on the active host and added to the standby host. This redirection method is applicable only when both the Master server and the Replica server are on the same IP subnet.

Choosing Off or On affects the available options in the Check With property for a Send ping request. Refer to Is Alive.

**Add IP/Mask**

Enter IPs for the active computer that will be moved to the standby computer during switchover. The Master IP address defined in the Master Properties must be different than the IPs entered here.

**To add IP/Mask**

1. Click the tree entry: Click here to add new IP/Mask.

   The IP Address dialog appears.

2. Enter the relevant IP/Mask data in the dialog and click OK.

A new entry is added to the list, and a new row opens for another IP/Mask. Enter as many entries as you require.

**Notes**:

- The Master IP address on the Properties tab under Host Connection (the Master IP that was entered in the Scenario Creation Wizard), must NOT be one of the IPs included in this list.

- If the Move IP or the Redirect DNS property is set to On, Arcserve RHA turns off the dynamic DNS registration for the Master. (The checkbox, Register this connection's addresses in DNS in the Advanced TCP/IP Settings dialog is cleared).

**Redirect DNS**

During the switchover, the A-record of the Master Server will be updated. This redirection option is applicable when the Master and the Replica are located on a different IP subnet or on the same subnet.

If the Move IP or the Redirect DNS property is set to On, Arcserve RHA turns off the dynamic DNS registration for the Master. (The checkbox, Register this

connection's addresses in DNS in the Advanced TCP/IP Settings dialog, is cleared).

**DNS Server IPs**

Enter IPs of DNS servers to update. Arcserve RHA tries to update all servers listed. However, switchover is considered successful even if only one update is successful.

To enter the value, click the tree entry: Click here to add new IP.

**DNS TTL**

Enter number of seconds for DNS Time-To-Live. This value is changed in the DNS Server for the A-record that is updated.

**Active Directory Integrated**

Specify if DNS is an Active Directory Integrated. If the Master DNS is on a Windows platform and integrated with Active Directory, set this option to On.

**DNS Key Filename (full path)**

Enter the full path of the file containing the DNS secure key. This field appears only when AD Integrated is Off.

**Master/Replica IPs in DNS**

Enter the Master/Replica IPs in its DNS server.

The Master's DNS record is updated during every switchover: in the switchover from Master to Replica, the addresses in the Master's DNS record are replaced by the addresses of the Replica. In the switch back from Replica to Master, the addresses are restored to the original values.

To enter the value, click the tree entry: Click here to add new IP.

**Switch Computer Name**

This redirection option is applicable when you use NetBIOS name resolution for your connections with the Master. If the hostname and the NetBIOS name are not the same, this option cannot be used.

During the switchover, the Replica computer is renamed to the Master computer name and the Master computer is renamed to a temporary name (if the Master server is alive). During the switchback, the names are restored. Both the hostname and the NetBIOS name are changed. When Switch Computer Name redirection is enabled in the Active Directory environment, the Master and Replica servers must be members of the same domain or trusted domain.

**Important!** If you will be redirecting File Shares, in which clients

connect to via the Master server name, Switch Computer Name must be enabled. For example, if the Master server's name is fs01 and clients connect to \\fs01\sharename or \\fs01.domain.com\sharename, then you must use the Switch Computer Name method. We also recommend enabling one other method. The most common method is to use both DNS Redirection and Switch Computer Name.

When you use the Switch Computer Name redirection method on Windows 2008 and Windows 2008 R2 systems, a reboot is required after switchover and switchback. You cannot run a reverse scenario until the system is restarted because the computer name does not take effect until the reboot completes. We recommend setting the Reboot After Switchoverand Switchback property to On when using this method.

### Master Computer Name

NetBIOS name of the Master computer. This name cannot be modified here.

### Replica Computer Name

NetBIOS name of the Replica computer. This name cannot be modified here.

### Reboot after Switchover and Switchback

After a switchover and a switchback, if this option is set to On, both Master and Replica computers are rebooted.

### Use Computer Name Alias

After a switchover, the Computer Name of the original Replica will not be changed. Instead, the Replica will use the Master hostname as the Alias, and therefore it will not have to reboot. If you set this option to On, we recommend you set the Redirect DNS property to On.

## User Defined Scripts

This option allows the standard redirection methods to be enhanced or replaced by actions that are invoked by user-defined scripts.

**Important!** When using scripts, each script must reside in the same path and with the same name on both the Master and the Replica.

### Active to Standby Redirection Script

- **Script Name**

Enter the name and full path of script to be run on the active computer, if it is alive, in order to redirect clients to the standby computer or release network resources on active host.

▪ **Arguments**

Arguments to be passed to the script specified in the previous property. Argument values must be static.

**Note**: The Active to Standby Redirection Script is also automatically executed when the HA scenario starts running. At this stage, the script runs on the standby Replica.

**Standby to Active Redirection Script**

▪ **Script Name**

Enter the name and full path of script to be run on the standby host, in order to redirect clients to it or add network resource.

▪ **Arguments**

Arguments to be passed to the script specified in the previous property. Argument values must be static.

**Identify Network Traffic Direction Script**

Required to fully support custom redirection methods. The custom scripts entered here are used to identify the active server. The Engine assumes that:

◆ If the script was executed on the host is returning 0, then the host is active (has all network resources on it or users directed to this host)

◆ If the script is returning a non-zero value, then the host is inactive (all or some of the network resources are absent, or users are not directed to this host).

▪ **Script Name**

Enter the name and full path of script to be run. This script determines if the Forward or Backward scenario will run when the scenario is started. The script runs on both Master and Replica: the one that returns zero is active. If both return the same value, a conflict is reported.

▪ **Arguments**

Arguments to be passed to the script specified in the previous property. Argument values must be static.

**Switch Computer Alias**

During switchover, the alias name is released from the active host and added to the standby host.

Both NetBIOS and DNS CNAME alias names are supported. To add NetBIOS alias names, create a Multi-String value named OptionalNames with Data as

<aliasnames> in the following registry key and restart the Server service.

HKEY_LOCAL_MACHINE\SYSTEM\Cur-
rentControlSet\services\LanmanServer\Parameters

**Note:** When the master server is a workgroup server, then the switchover by DNS alias name cannot be supported. The reason is that the user of the workgroup server does not have the privileges to access the DNS information.

### Alias Name

Enter the alias name of the active host that is moved to the standby host dur-ing switchover.

### Update/Add/Delete

By default, alias name is auto configured when the HA scenario is created. You can add or delete all or the selected aliases.

After the scenario is created, you can also add a new alias in your envir-onment. To enable this new alias, click the refresh button and reload the HA properties.

# Is Alive

Arcserve RHA continuously checks to see if the active host is up (according to *Send pingrequest*, *Connect to DB*, *or User-Defined Script* method*, see below). These checks are made in scheduled intervals according to the Heartbeat Frequency.

Checking the status is interpreted as follows:

- If there is an indication that the active computer is OK (alive), no new action is taken, and the system continues to check according to the scheduled intervals of the Heartbeat Frequency.

- If there is an indication that the active computer is not OK (is not alive), the active computer is checked again at the next Heartbeat timeout for a maximum period of the Is Alive timeout. If no indication of the active host being alive is found within the Is Alive timeout, Arcserve RHA executes event notification. Simultaneously, it checks whether or not to perform a switchover, as defined by the Perform switchover automatically property.

**Important!** When using scripts, each script must reside in the same path and with the same name on both the Master and the Replica.

**Is Alive Timeout (sec)**

If the standby host does not receive indication that the active host is alive during this interval (in seconds), switchover or notification is initiated. The checks are performed at the Heartbeat Frequency.

Default is 300 seconds.

**Heartbeat Frequency (sec)**

Interval (in seconds) for sending heartbeat requests (performing the checks below).

Default is 30 seconds.

**Check Method**

**Send Ping Request**

ICMP requests are sent automatically from the standby host to the active host, to check if the active host is up.

The available options depend on the value of the Move IP property. For more information, refer to Redirection properties.

- If **Move IP** is **On**

- During the switchover, the IP is moved from the active computer to the standby. Therefore, the standby computer must check this IP continuously.

  In the **IP for Ping** property, enter IP Address to ping.

- If **Move IP** is **Off**

  During the switchover, the IP is not moved from the active to the standby computer. Therefore, define two IPs for ping:

  **IP for ping from Master to Replica**
  Enter IP address to ping. When the Replica computer is the active host, an ICMP request is made from the Master to the Replica. If no reply is received within two seconds, the Replica computer is considered non-operational.

  **IP for ping from Replica to Master**
  Enter IP address to send ping to. When the Master computer is the active host, an ICMP request is made from the Replica to the Master. If no reply is received within 2 seconds, then the Master computer is considered to be non-operational.

**Connect to DB**

[For database applications only] When this property is set to On, Arcserve RHA connects to the active computer's database according to the Heartbeat Frequency, and checks to see if the database services are running and all databases are mounted.

**User-Defined Scripts**

Allows the standard check methods to be augmented or replaced by user-defined actions in scripts.

**Check Script on Active Host**

This script runs on the active server, and checks if it is alive.

- **Script Name**

  Enter the name and full path of script to run. Arcserve RHA connects to the active computer once every Heartbeat timeout, and executes the script. If the return value equals zero, the active computer is OK (alive). If

the return value is different than zero, the active server is not responding and switchover is required.

◆ **Arguments**

Arguments to be passed to the script specified in the previous property. Argument values must be static.

**Check script on Standby host**

This script runs on the standby server, and checks if the active server it is alive.

◆ **Script Name**

Enter name and full path of script to be run.Arcserve RHA connects to the active computer once every Heartbeat timeout, and executes the script. If the return value equals zero, the active computer is OK (alive). If the return value is different than zero, the active server is not responding and switchover is required.

◆ **Arguments**

Arguments to be passed to the script specified in the previous property. Argument values must be static.

# DB Management/Application/Shares Management

**Automatic**

If you want Arcserve RHA to manage services on your DB Server, set this option to On. Then:

1. During the scenario creation, the list of application (DB) services that must be managed are auto-discovered.

2. Once the scenario is running:

   ◆ [DB] DB services on the active host are initiated (if they are not running), and they are stopped on the standby host (if they are running).

3. During the switchover:

   ◆ [DB] DB services on the active host are stopped, and they are started on the standby host.

**User-Defined Scripts**

**Start DB/Start Application Script**

If set to On, runs a user-defined script to augment or replace the start of DB services/the start of applications. This action occurs during a scenario run on the active host, or during a switchover on the standby host.

**Script Name (full path)**

Enter the name and full path of the script to be run.

▪ **Arguments**

Arguments to be passed to the script specified in the previous property. Argument values must be static.

**Stop DB/Stop Application Script**

If set to On, runs a user-defined script to augment or replace the stop of DB services/the stop of applications. This action occurs during a scenario run on the standby host, or during a switchover on the active host.

▪ **Script Name (full path)**

Enter the name and full path of the script to be run.

▪ **Arguments**

Arguments to be passed to the script specified in the previous property. Argument values must be static.

# Actions upon Success

**Important!** When using scripts, each script must reside in the same path and with the same name on both the Master and the Replica.

**User-Defined Script**

When set to On, runs a user-defined script. The actions invoked by the script will be performed following the completion of a successful switchover.

- **Script Name (full path)**

Enter the name and full path of script. This script runs on the active server after the switchover completion.

- **Arguments**

Arguments to be passed to the script specified in the previous property. Argument values are static.

# Active and Standby Hosts

In an initial scenario, the Master is the active computer, and the Replica is the standby computer. The standby computer is continuously checking the state of the active one, to decide whether to become active computer.

The first time a switchover occurs, the Replica that was on standby becomes the active computer, and the Master reverts to a standby mode (assuming it is still operational). When the Master (now the 'standby') is ready, a switchback process can be initiated where the Master again becomes active, and the Replica returns to its previous standby and monitoring role.

# Move IP Redirection

This section describes the steps required for adding Move IP redirection to the High Availability scenario.

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

This section contains the following topics:

- Adding IP Address to the Master Host
- Configuring the Move IP Method through the Manager
- Cluster Move IP

# Adding IP Address to the Master Host

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

**To add IP address to the Master host**

1. Select **Start**, **Settings**, **Control Panel**, **Network Connections**, **Local Area Connection**.

   The **Local Area Connection Status** dialog appears.

2. On the **General** tab, click the **Properties** button.

   The **Local Area Connection Properties** dialog appears.

3. On the **General** tab, select **Internet Protocol (TCP/IP)**, and then click the **Properties** button.

   The **Internet Protocol (TCP/IP) Properties** dialog appears.

4. On the **General** tab, click the **Advanced** button.

   The **Advanced TCP/IP Settings** dialog appears.

5. On the **IP Settings** tab, click the **Add** button.

6. The **TCP/IP Address** dialog appears.

7. In the **TCP/IP Address** dialog, enter the additional IP address (RHA-IP). Then, click Add.

   The additional IP address is saved, and the **TCP/IP Address** dialog is closed.

8. Click **OK** on all the open dialogs, until you close all dialogs and exit the local area connection settings.

# Configuring the Move IP Method through the Manager

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

- For new scenarios, directly from the Scenario Creation Wizard.
- For existing scenarios, by modifying the Master host name.

The procedures for both ways follow.

This section contains the following topics:

- Adding RHA-IP to New Scenarios
- Adding RHA-IP to Existing Scenarios

# Adding RHA-IP to New Scenarios

**To add RHA-IP to a new scenario for Move IP redirection method**

1. On the Scenario Creation Wizard, in the **Master and Replica Hosts** page, enter the following:

   - On the **Master Hostname/IP** box, enter the additional IP address (RHA-IP).

   - On the **Replica Hostname/IP** box, enter the IP address of the Replica host, and not its hostname.

2. Click **Next**, and continue defining the scenario as usual until the **Switchover Properties** page appears.

3. On the **Switchover Properties** page, open the **Network Traffic Redirection** group, select the **Move IP** property, and set its value to On.

   By default, the second IP address of the Master host appears here in the **IP/Mask** box.

   **Note**: If the Master host has only one IP address, the **IP/Mask** box would be empty.

4. If you have end users who connect to the Master host using its hostname, use the **Redirect DNS** or **Switch Computer Name** methods along with the **Move IP**. If you do not need to use the Master hostname, disable the **Redirect DNS** option by setting its value to Off.

5. After setting the redirection method, click **Next**, and continue defining the HA scenario as usual.

# Adding RHA-IP to Existing Scenarios

To perform the following, stop the scenario first.

**To add RHA-IP to an existing scenario for Move IP redirection method**

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

1. Right-click the Master and select **Rename** from the pop-up menu. Then, enter the **RHA-IP** address.

2. Make sure that the Replica host is defined by its IP address and not by its hostname. If necessary, enter the Replica IP address instead of its hostname.

3. On the Framework pane, select the **High Availability Properties** tab..

4. Open the **Network Traffic Redirection** group, select the **Move IP** option, and set its value to On.

   The **IP/Mask** property appears.

5. Click the **IP/Mask** value box.The **IP Address** dialog appears.

6. Enter the original IP address of the Master host. This IP address will be moved to the standby computer during switchover. Then, click **OK**.

   **Note**: 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.**

7. If you have end users who connect to the Master host using its hostname, use the **Redirect DNS** or **Switch Computer Name** methods along with the **Move IP**. If you do not need to use the Master hostname, disable the **Redirect DNS** option by setting its value to Off.

8. Click the **Save** button on the Standard toolbar to save your setting.

# Cluster Move IP

Using Move IP redirection with a clustered Master (MSCS with shared storage) requires you to add an IP resource to the Master resource group. This section describes how to configure this redirection method.

**Note:** If both Master AND Replica are clusters, perform the following steps:

1. Manually create an IP resource with the IP that you want to move to the replica cluster and make the resource offline.

2. Create an HA scenario as usual and use the Move IP redirection method. Make sure that the IP resource you created on the replica cluster is the same IP that you want to move.

3. Run the scenario as usual.

This section contains the following topics:

- Using the Manager
- Using the Master Cluster

# Using the Manager

This section details Cluster Move IP redirection using the Manager.

- For New Scenarios
- For Existing Scenarios

# For New Scenarios

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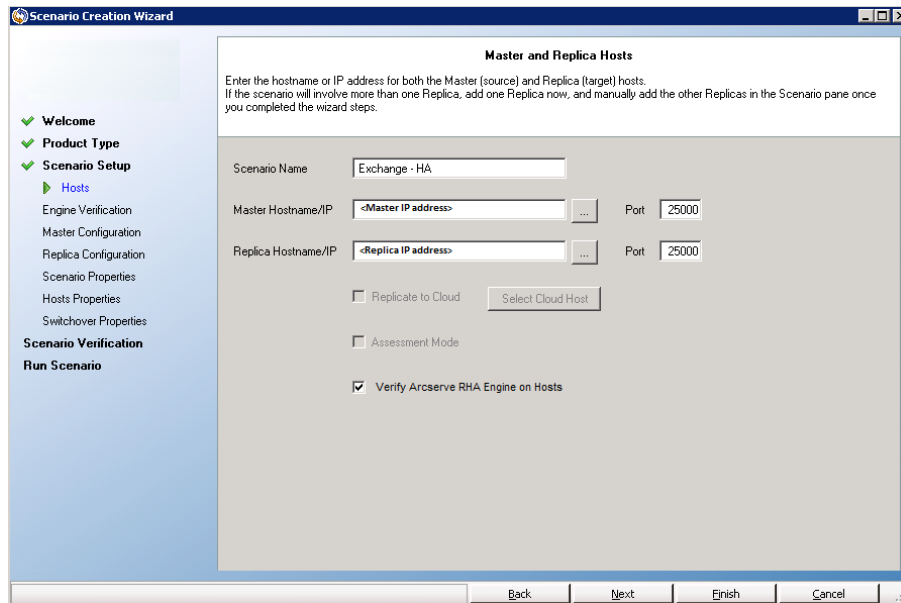
# For Existing Scenarios

**To use Cluster Move IP with existing scenarios**

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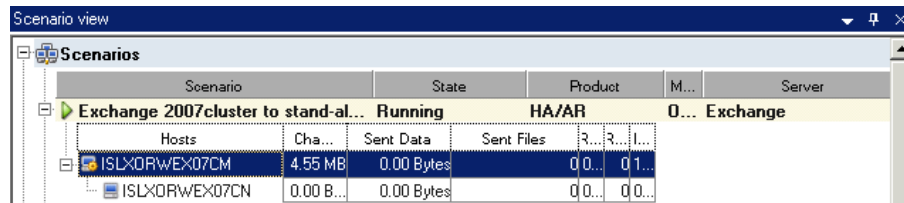


2.  Right-click the Master and select **Rename** from the pop-up menu. Then, enter the RHA-IP address.

3.  On the Framework pane, select the High Availability Properties tab and then select the Replica server as the switchover host.

4.  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 mul-

tiple production IP addresses by selecting **Click here to add new IP/Mask.**

# Using the Master Cluster

**To use Cluster Move IP through the Master cluster**

1.  Open the Cluster Administrator.

2.  In the Master Cluster Resource Group, create a new IP resource and name it **RHA-IP**.

3.  Bring this resource online, and verify 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.

# Chapter 11: Protecting the Control Service

This section explains how to replicate the Control Service data, and how to switch the roles of two Control Services when the active Control Service is down. The section describes in detail the following operations: creating and using HA Control Service scenario, and performing Control Service switchover and switchback.

This section contains the following topics:

# Understanding the Control Service Scenario

The Control Service functions as the single-point-of-control of the Arcserve RHA operation, and it contains the entire data of the existing scenarios. In one Arcserve RHA scenario-system, one Control Service manages all scenario-related-tasks, and the Managers that are connected to it enable you to monitor the Arcserve RHA activities. If the Control Service is down, the scenario functioning is not affected. However, you are not able to control, manage and monitor the state and operation of the scenarios during this time. To overcome the danger of losing the Control Service data or losing the ability to manage and monitor your scenarios, Arcserve RHA offers you the Replication and HA Control Service scenarios. These scenarios enable you to protect the Control Service data and functionality, in the same way you protect other supported applications.

Arcserve RHA enables you to replicate the Control Service data, and to save the replicated data on a Replica host. In order to perform this, you need to create a Replication Control Service scenario. The Replication scenario also enables you to activate the Rewind option, and to recover lost Control Service data if necessary.

In addition, Arcserve RHA enables you to apply the HA solution to the Control Service. This means that if the active Control Service is down, you can switch the roles between the active Control Service and a standby Control Service, and make the standby Control Service the active one. For switching over and switching back the roles of two Control Services, you need to create an HA Control Service scenario.

**Important!** If you are running a Replication Control Service scenario, you cannot use a second Control Service to manage your scenario-related-tasks. To use a second Control Service when the first one is down, you need to initially install two Control Services, one as the active Control Service and the second as the standby Control Service. You also need to install two Engines, one on each Control Service machine, and to verify that they are running. Only then you can create and run HA Control Service scenario.

The creation of Replication and HA scenarios for Arcserve RHA Control Service is similar to the creation of Replication and HA scenarios for application and database servers. In both you are using the same step-by-step Scenario Creation wizard. However, there are some differences in the creation of Replication and HA scenarios for the Arcserve RHA Control Service, as follows:

- [Replication and HA scenarios] Running only one scenario per Control Service - you can run only one Control Service scenario at a time for a specific Control Service.

- [Replication and HA scenarios] No special license is needed - you do not need a special license for creating a Control Service scenario, either Replication or HA. However, you do need to register the Arcserve RHA product before creating a scenario for the Control Service.

- [Replication and HA scenarios] Master details cannot be changed - In the **Master and Replica Hosts** page in the Scenario Creation Wizard, where you enter the IP address/hostname of the Master and Replica hosts, the Master host details are entered automatically by the system and cannot be changed. The Master Control Service details that appear in the wizard are the ones you entered in the Web browser for connecting the Control Service to the Overview Page.

- [HA scenario] Control Service items cannot be excluded from replication - in the **Master Configuration** page in the Scenario Creation Wizard, the auto-discovery results are read-only. You cannot exclude Control Service items from the replication process.

- [HA scenario] Move IP redirection method cannot be used - there are only two network traffic redirection methods you can use: Redirect DNS and Switch Computer Name. You cannot use the Move IP redirection method.

- [HA scenario] Automatic switchover and automatic reverse replication cannot be disabled - you cannot disable the automatic initiation of a switchover when the Master is down, and the automatic initiation of a backward scenario. Therefore, the **Switchover and Reverse Replication Initiation** page in the Wizard and the corresponding properties are either not displayed or disabled. However, you can manually initiate a switchover and a switchback by using the **Perform Switchover** button on the Standard toolbar.

- [HA scenario] To create HA scenario for a Control Service, you need to install two Control Services: one should function as the active Control Service, and the other should function as the standby Control Service. For more information, refer to *Arcserve RHA Installation Guide*.

To learn how:

- To create Replication Control Service scenario, use the instructions for Create a File Server Replication Scenario, along with the qualifications specified above.

- To recover Control Service data, refer to the Recovering Data and Servers chapter.

- To create HA Control Service scenario, refer to Create High Availability Scenarios for the Control Service.

- To manually initiate a switchover, refer to Manually Initiating a Control Service Switchover.

- To handle a loss of connection and the switchover process, refer to The Switchover and Backward Scenario Processes.

- To reverse back the Control Services o their original states, refer to Switching Back the Control Services Roles.

# Create High Availability Scenarios for the Control Service

When creating Control Service scenarios, Assured Recovery is not supported and the option is unavailable.

**Important!** Before you create HA scenario for the Control Service, you should verify that you have two Control Services installed, one as the (active) Master Control Service and one as the (standby) Replica Control Service. The Replica Control Service should be down. In addition, an Engine should be installed and running on both the Master and Replica hosts.

**To create a High Availability scenario for the Control Service**

1.  Open the Arcserve RHA Manager. Then, select from the **Scenario** menu the

    **New** option, or click the **New** button on the Standard toolbar.

    The **Scenario Creation Wizard** opens.

2.  Select the required scenario options, as follows:

    ◆ Select the **Create a New Scenario** option button.

    ◆ From the **Group** drop-down list, select the group to which you want to assign the new scenario, or enter a name for a new scenario group.

3.  Click **Next**. The **Select Server and Product Type** page opens.

    A list of available applications and scenario types is presented.

    **Note:** The list of available applications depends on the licenses applied**.**

    Select the required scenario options, as follows:

    ◆ From the **Select Server Type** list, select **Control Service**.

    ◆ From the **Select Product Type** options, select **High Availability Scenario (HA)**.

    ◆ Note that Integrity Testing for Assured Recovery is not supported for Control Service HA scenarios.

4.  Click **Next**. The **Master and Replica Hosts** page opens.

5.  Enter the following information:

    ◆ In the **Scenario Name** box - accept the default name or enter a new name for the scenario. When entering a name, choose a unique name, since you cannot use the same name for more than one scenario.

- ◆ In the **Master Hostname/IP** box - the system automatically enters the hostname or IP address of the (active) Master Control Service, based on the Control Service details you entered for opening the Overview Page. These details cannot be changed here. To use a different Control Service, you need to connect it to the Overview Page, and then reopen the Manager.

- ◆ In the **Replica Hostname/IP** box - enter the hostname or IP address of the Replica (standby) Control Service, or use the **Browse** buttons to find it.

- ◆ In the **Port** boxes - accept the default port no. (25000) or enter a new port numbers for the Master and Replica.

   **Notes:**

   - ◆ The **Assessment Mode** option is disabled in HA scenario.

   - ◆ The **Verify Arcserve RHA Engine on Hosts** check box - select this check box if you want the system to verify whether Engines are installed and running on the Master and Replica hosts you specified in this page. If Engines are not installed on the selected hosts, you can use this option to remotely install the Engines on one or both hosts. For more information about the **Host Verification** page, see [Creating a Scenario](#).

6. After you selected the desired options, click **Next**. The **Master Configuration** page opens.

Arcserve RHA auto-discovery component automatically displays the directories and files that are on the active Control Service. These directories and files are the data that will be replicated and protected.

The replicated Control Service items include:

- ◆ Product Registration - product registry keys

- ◆ Scenarios - xmc files of scenario definitions

- ◆ Templates - xmc files of user-defined templates

- ◆ Reports - files of scenario reports

- ◆ Configuration and Management files

**Note**: In HA Control Service scenarios, you cannot exclude Control Service items from replication. In Control Service Replication scenarios, you can exclude the items you do not want to replicate, by clearing their check boxes.

7. Click **Next**. The **Scenario Properties** page opens.

The **Scenario Properties** page enables you to configure the scenario properties that affect the entire scenario. Typically, the default values are sufficient.

If you want to configure the scenario properties at this stage, refer to Understanding Scenario Properties. To configure the scenario properties at a later stage, refer to Configuring Scenario Properties.

8. Click **Next**. The **Master and ReplicaProperties** page opens.

The **Master and Replica Properties** page enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient.

If you want to configure the Master and Replica properties at this stage, refer to Setting Master and Replica Properties. To configure the Master and Replica properties at a later stage, refer to Configuring Master or Replica Server Properties.

**Note**: You can modify all the settings in this pane after the scenario is created. However, before changing any Spool properties (which can be configured here), review the Spool information for configuration details.

9. Once you are satisfied with the Master and Replica properties, click **Next**.

The **Switchover Properties** page opens.

The **Switchover Properties** page allows you to modify switchover parameters. As with the prior steps, no changes are required.

If you want to configure the switchover properties at this stage, refer to Understanding High Availability Properties. To configure the switchover properties at a later stage, refer to Configuring High Availability Properties.

**Notes:**

◆ When selecting the **Network Traffic Redirection** method, there are only two methods you can use for this scenario: **Redirect DNS** and **Switch Computer Name**. You cannot use the **Move IP** redirection method.

◆ The **Is Alive Timeout (sec)** property controls how long to wait after a failure is detected before triggering an automatic switchover. The default is 300 seconds. For more details, review the Is Alive information.

10. Click **Next.** A notification message appears informing you that <caha> verifies the validity of the new scenario and checks many different parameters between the Master and Replica servers to ensure a successful switchover.

**Note**: In HA Control Service scenario, once a Master failure is detected, a switchover and a backward scenario are always initiated automatically. You

cannot disable this automatic initiation. However, you can also manually initiate a switchover, by clicking the **Perform Switchover** button.

11. Once the verification is completed the **Scenario Verification** page opens.

12. If the scenario was not set up correctly, or problems occurred in the participating hosts or the connection between the Arcserve RHA components, the errors and warnings detected are displayed, and two additional buttons appear: **Retry** and **More Info**.

13. To repeat the verification process, click the **Retry** button.

14. To view additional information about the errors and warnings, click the **More Info** button.

   The **Verification Results** dialog opens, listing all the errors and warnings detected.

   The **Verification Results** dialog provides you with detailed information about the checks performed to help diagnose problems. It is intended to help you resolve any issues encountered in running the software. You can also contact Technical Support for further assistance.

   - If any errors are displayed, you cannot run the scenario. These errors must be corrected before you can start the synchronization, replication and HA processes.

   - If only warnings are displayed, you can run the scenario. However, it is important that you consider the warning carefully since they indicate conditions that are known to potentially cause problems with replication or switchover. To close the dialog and return to the **ScenarioVerification** page, click the **Cancel** button.

   - When the scenario is verified successfully, on the **Scenario Verification** page click **Next** to continue.

   The **Scenario Run** page opens.

15. The scenario configuration is now completed and you are prompted to run it. Running the scenario starts the data synchronization process, following by replication and is alive checks.

   - To finish the scenario creation and run it later, select **Finish**.

   - To run the scenario, click **Run Now**.

16. The synchronization process starts. Synchronization may take a while depending on data size and network bandwidth between the Master and Replica. You

will receive the following message in the Event pane when synchronization is complete: **All modifications during synchronization period are replicated**.

From this point, real-time replication occurs and the High Availability solution is active.

# Open the Manager for Using the HA Control Service Scenario

To properly work with the HA Control Service scenario, it is important that you open the Overview Page, and from it the Manager, by using the Control Service hostname, instead of its IP address. If you will use the Control Service IP address, after a switchover Arcserve RHA will not be able to automatically reconnect the Overview Page and the Manager to the new active Control Service.

In addition, if you intend to work with HA Control Service scenario, you should NOT open the Manager from the machine where you installed a Control Service. Open the Manager from a third machine, which does not act as either the active or standby Control Service.

**To open Arcserve RHA Manager for working with HA Control Service scenario**

1. Open Internet Explorer. On the **Address** box, enter the Control Service host-name and Port Number as follows: http://host_name:port_no/start_page.aspx

   **Note**: If you selected the **SSL Configuration** option during the installation of the Control Service**,** enter the Control Service hostname and Port Number as follows: https://host_name:port_no/start_page.aspx

   The **Login** dialog opens.

2. Enter your User Name, Password and Domain and click the **Log In** button.

   The **Overview page** opens.

3. On the **Quick Start** toolbar on left, click the **Scenario Management** option.

   A progress bar appears, indicating that the Manager component is currently installed on the local machine.

4. Once the Manager installation is completed, the Manager opens.

   Now, you can now start creating the HA Control Service scenario.

# Switch the Roles of the Active and Standby Control Services

The shutting down of a Control Service, either intentionally or due to a crash, presents a unique problem. When the Control Service is down, the Overview Page and the Manager are disconnected. Consequently, they can no longer receive updated information, and they cannot display a visible indication to the type of event that occurred. Although the Overview Page and the Manager are up, you cannot see that the Control Service is down, and you cannot manually initiate a switchover at this stage.

Arcserve RHA handles the disconnection problem by automatically trying to restore the Control Service ("Manager") to its active state. If the attempt fails, and the active Control Service is still detected as down, Arcserve RHA automatically initiates a switchover. During the switchover, the standby Control Service becomes the active Control Service. Following that, the Overview Page and the Manager are automatically reconnected to the new active Control Service, and once again they display the updated state of your system. During the reconnection, you may be prompted to log in again.

When the original active Control Service is up again, Arcserve RHA automatically initiates a backward scenario. The backward scenario is a replication in the reverse direction: from the new active Control Service server to the new standby Control Service server. At this stage, you can reverse back the roles of the Control Services. All you need to do is to manually initiate a switchback, meaning, a switchover in the opposite direction.

In addition to the default automatic switchover, you can also manually initiate a switchover between the active and standby Control Services. Once triggered, whether manually or automatically, the switchover process itself is fully automated.

There are several stages in the role switching of the active and standby Control Services:

1. Initiating a switchover - this can be done either automatically by the system, when it detects that the active Control Service is down, or manually by you.

2. The switchover process and the automatic initiation of a backward scenario - these processes are performed automatically and you cannot disabled them.

3. Initiating a switchback - this can only be done manually by you, when you decide that the original active Control Service can become the active server again.

# Manually Initiating a Control Service Switchover

When Arcserve RHA detects that the active Control Service is down, it automatically tries to restart the Control Service, and if this attempt fails, it initiates a switchover. However, you can also manually initiate a switchover, when the active Control Service is still up.

**Note**: Do not run the Control Service on both the Master and Replica hosts at the same time to prevent unexpected errors, such as 'connection broken'.

**To manually initiate a switchover**

1. Open the Manager and select the desired Control Service 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 opens.

3. Click **Yes** on the **Perform Switchover** confirmation message. This procedure initiates a switchover from the active Control Service to the standby Control Service.

   From this stage on, the switchover process is the same for both manual and automatic initiation.

# The Switchover and Backward Scenario Processes

**Understanding the switchover and backward scenario process**

1. Since the original active Control Service is down, the Overview Page and Manager are no longer connected to it. Therefore, they no longer receive and display updated information, and the changes that occur following the switchover initiation are not shown in them, as they are shown in a regular switchover.

2. When the Overview Page loses its connection to the original active Control Service, the following message opens.

   This message indicates that the original active Control Service is down, and therefore it is no longer connected to the Overview Page.

3. Click **OK** to close the message. This message may appear several times until the original standby Control Service becomes active, and a connection to it is established.

4. When the original standby Control Service is up and functioning as the new active Control Service, the Overview Page is automatically reconnected to it, and the **Login** dialog appears, prompting you to login to the new active Control Service.

5. Enter your User Name, Password and Domain and click the **Log In** button.

   The **Overview page** re-appears, and it is now connected to the new active Control Service.

6. On the Manager, the **User Credentials** dialog may open.

   The **User credentials** dialog prompts you to login to the new active Control Service. If this dialog appears, enter the necessary details and click **OK**.

   **Note**: The appearance of the **User credentials** dialog is related to internal caching settings, and it is not necessarily an indication to the progress of the switchover process. The switchover may take place even if the **User credentials** dialog does not appear.

7. The two Control Services have switched roles. Now, the Manager is no longer connected to the original active Control Service but to the standby Control Service, which became active following the switchover. The switchover related-events are displayed in the Event pane.

   **Note:** The "Split Brain" problem and solution:

   After a connection loss and a switchover, the original standby Control Service is functioning as the active Control Service. However, the original active Control

Service may still be up. Upon reconnection, both Control Services may try to act as the active Control Service. In order to solve this potential problem, Arcserve RHA keeps built-in numerical parameter in each Control Service, and the switchover process increases the number of the newly active Control Service. All connection requests are sent with this parameter, and when a Control Service receives a connection request, it checks whether it contains a lower or higher number than the one it carries. The Control Service that carries the lower number, shuts itself down, and becomes the standby Control Service.

8. After the switchover, a backward scenario is automatically initiated by the system.

9. The backward scenario starts running once the original active Control Service is up. It replicates data from the new active Control Service to the new standby Control Service, while overwriting the data on the current standby Control Service.

10. Now, you can switch back the roles of the active and standby Control Service, and make the original Master the active server again and the Replica the standby server.

# Switching Back the Control Service Roles

When the original active Control Service is up again, and the backward scenario is running, you can switch back the roles of the standby and active Control Service and reverse them back to their original states.

**To initiate a switchback**

1. To reverse back the roles of the Control Services, while the backward scenario is running, click the **Perform Switchover** button or select the **Perform Switchover** option from the **Tools** menu.

   A confirmation message opens.

2. Click **Yes** on the **Perform Switchover** confirmation dialog. This procedure initiates a switchback from the original Replica server to the Master server. Again, you are not able to see the process of the switchback on the Overview Page and the Manager, since they are disconnected from the active Control Service. But once the Overview Page and Manager are reconnected, you can see that the Control Services have switched their roles and returned to their original states.

   Now, the Control Service HA scenario is running in its original direction.

# Chapter 12: Assured Recovery Testing

This section explains the Assured Recovery testing option, and describes the following operations: creating AR scenario, performing AR test in a scheduled and non-scheduled mode, and configuring the AR properties. In addition, this section describes how to set up VSS snapshot creation and how to manage snapshots.

This section contains the following topics:

# About Assured Recovery

The Assured Recovery option enables you to perform a full transparent test of the recoverability of your data on the Replica server. The Replica server that is tested is the one that would take over the production server if it will be down. The Assured Recovery option is a true test of the actual server, applications and actions that will be required in the event the Replica server will have to switch, become the Active server, and carry out its functions.

This Assured Recovery test is executed by starting up database services, and performing whatever operations that are required to verify the integrity of the data. All this is done without any need to perform resynchronization, and without impacting either the availability of the production server, or the safety that the Replication and HA systems are designed to provide.

During the test, the data changes that continue to take place on the Master are sent to the Replica, but they are not immediately applied. Instead, these changes are accumulated and stored in a spool, and only when the testing is completed they are applied to the Replica data. Since the spooling occurs on the Replica, if something happens to the Master during the testing process, none of the accumulated changes are lost.

Once the testing is finished, the Assured Recovery option stops the application services it started on the Replica. Then, the Replica server is automatically rewound to precisely the state that existed when the replication was paused and the test started. This way, the accumulated changes in the spool can be applied as if no testing has occurred. From this point on, the Replication or HA scenario continues normally. In the case of an HA scenario, if a failure of the Master occurred during the testing, switchover begins.

The Assured Recovery test can be fully automated and performed on a scheduled basis as often as needed. Upon completion, appropriate personnel can be alerted with the status of the test, and additional actions can be triggered on success, for example, taking a VSS snapshot of the Replica data or creating a backup. In addition, you can perform AR testing in non-scheduled mode when the need arises.

The Assured Recovery testing is tailored to all supported application and database servers. However, since the Assured Recovery option tests database services, it is less applicable for File and IIS Servers. You can still use the Assured Recovery option with these servers for special tasks. For example, you can automatically suspend replication on a regular basis during several hours each day, week or month, and run scripts in this interval, or you can use this suspension to take VSS snapshots on

the Replica. Since there is no *application* per se, testing of the data with File and IIS Servers scenarios requires additional custom scripts.

The Assured Recovery option supports both Replication and HA solutions, except for Control Service scenarios. However, it is best suited for HA since in this case the Replica server necessarily contains the actual database servers, on which the test is performed, and not only data.

**Note**: The Assured Recovery option is not available for Control Service scenarios.

If you are using AR test as a part of Replication scenario, you must verify that the root directories path is the same on the Master and the Replica. In addition, the Replica should have database application installed, or share files if you test a File Server, and they need to be configured on the Master and the Replica in exactly the same way. Otherwise, the AR test will not produce meaningful results.

# Creating Assured Recovery Testing Scenarios

The Assured Recovery testing feature must be enabled during the creation of the scenario that will later use it. For this reason, you cannot perform testing within a Replication or HA scenario that is already running, and was not configured to use the Assured Recovery option. To use Assured Recovery, it is necessary to create a new scenario with the Integrity Testing for Assured Recovery option turned to On.

**Note**: This section demonstrates the creation of an Assured Recovery testing scenario for Exchange Server HA. The procedure is similar for all application types.

**To set Assured Recovery testing scenario**

1. Open the Arcserve RHA Manager. Then, select from the Scenario menu the New option, or click the New button on the Standard toolbar.

   The Scenario Creation Wizard opens.

2. Select the required scenario options, as follows:

   ◆ Select the Create a New Scenario option button.

   ◆ From the Group drop-down list, select the group to which you want to assign the new scenario, or enter a name for a new scenario group.

3. Click Next. The Select Server and Product Type screen opens.

4. A list of available applications and scenario types is presented.

   **Note:** The list of available applications depends on the licenses applied**.**

   Select the required scenario options, as follows:

   ◆ From the Select Server Type list, select the type of server for which you want to create the AR scenario. For this example, we will use Microsoft Exchange Server.

   ◆ From the Select Product Type options, select either Replication and Disaster Recoveryor High Availability Scenario.

   **Note**: The Assured Recovery test is best suited for HA scenarios. If you select the Replication option, you must verify that the root directories path is the same on the Master and the Replica. In addition, the Replica should have database application installed, or share files if you test a File Server. Otherwise, the AR test will NOT produce meaningful results.

   ◆ Select the Integrity Testing for Assured Recovery option.

5. Click **Next**. The Master and Replica Hosts screen opens.

6. Enter the following information:

   ◆ In the Scenario Name box - accept the default name or enter a new name for the scenario. When entering a name, choose a unique name, since you cannot use the same name for more than one scenario.

   ◆ In the Master and Replica Hostname/IP boxes - enter the hostname or IP address of the Master (active) and Replica (standby) servers, or use the Browse buttons to find them.

   **Important!** Only one Replica can be configured for AR testing in a single scenario. If, at a later stage, you will add a Replica to the scenario and try to configure it for the AR test, the following message will appear: **Only one scheduled task per scenario can be set. Replica integrity testing for Assured Recovery for host [Replica_name] is already switched on. Do you want to turnthis option off now?** To switch the test to the second Replica, you will need to click Yes.

   **Note:** If either server is a MSCS cluster, enter the Virtual Server Name or IP address as the Master and/or Replica name (instead of the physical node's name/IP).

   ◆ In the Port boxes: accept the default port no. (25000), or enter a new port numbers for the Master and Replica.

   ◆ The Verify Engine on Hosts option - select this check box if you want the system to verify whether Engines are installed and running on the Master and Replica hosts you specified in this screen. If Engines are not installed on the selected hosts, you can use this option to remotely install the Engines on one or both hosts.

7. After you entered or selected the desired options, click Next. The Databases for Replication screen opens.

   The auto-discovery component automatically displays the Exchange databases that are on the Master server. These are the databases that can be replicated and protected.

8. By default, all the discovered databases are selected and all will be replicated. You can exclude any of these storage groups from replication by clearing their check boxes.

9. Click Next. The Replica Configuration screen opens.

   The auto-configuration component verifies that the Exchange Server configuration on the Master and Replica servers will be identical during the replication procedure. This means that if there are discrepancies, Arcserve RHA

will perform the required actions, including: deleting storage groups, public folders or mailbox stores from the Replica, create new ones and make modifications to existing ones. The actions that will be performed during the configuration process are indicated in the Action column on the right.

10. Review the changes that will occur during the automatic configuration on the Replica Exchange server, and make sure you want them to be performed.

    **Note:** If a Remove action is indicated, make sure that you are ready to delete the specified storage item from the Replica server, since it does not have an automatic backup. If you want to save it in a different location before deletion, click the Finish button to exit the wizard.

    **Important!** You cannot use UNC paths as root directories on the Replica host for Assured Recovery scenario.

11. Click **Next** to start the Replica configuration process. The Scenario Properties screen opens.

    The **Scenario Properties** screen enables you to configure the scenario properties that affect the entire scenario. Typically, the default values are sufficient.

    If you want to configure the scenario properties at this stage, refer to Understanding Scenario Properties. To configure the scenario properties at a later stage, refer to Configuring Scenario Properties.

12. Click Next. The Master and Replica Properties screen opens.

    The Master and Replica Properties screen enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient.

13. To verify that the Assured Recovery option is active, under the Replica Properties list on the right, open the Scheduled Tasks group and ensure that the Replica Integrity Testing for Assured Recovery property is set to On. You can leave the default values of the other related properties, and change them later if needed. For more information about AR properties refer to Understanding Assured Recovery Properties.

    If you want to configure the Master and Replica properties at this stage, refer to Setting Master and Replica Properties. To configure the Master and Replica properties at a later stage, refer to Configuring Master or Replica Server Properties.

    **Note**: You can modify all the settings in this pane after the scenario is created. However, before changing any Spool properties (which can be configured here), review the Spool information for configuration details.

14. Click Next. If you selected HA, the Switchover Properties screen opens.

15. From this stage, set up the scenario as you would normally following the instructions in the usual manner. For more information, see the appropriate Operation Guide. After the scenario creation is completed, run the scenario.

Once the initial synchronization is completed and the replication process is active, the AR test can be performed.

# Configure Assured Recovery Properties

To configure Assured Recovery properties, the scenario must be stopped.

**Note:** The Properties pane and its tabs (Root Directories, Properties, Statistics) are context sensitive, and change whenever you select a different node from a scenario folder.

**To set Assured Recovery scenario properties**

1. On the Scenario pane, select the Replica that you want to test and whose properties you want to configure.

2. On the Framework pane, select the Properties tab.

   The Replica Properties list opens.

   **Note**: A running scenario has a gray background, and scenarios that are not running have a white background.

3. If the scenario is running, click the Stop button on the toolbar. The scenario is stopped.

4. On the Replica Properties list, open the Scheduled Tasks group to display the Replica Integrity testing for Assured Recovery properties.

5. From the list, select the required property, and select or enter the appropriate values. Some values can be selected from a combo box while other values can be manually entered in an edit box field.

6. Click the Save button on the Standard toolbar to save and apply your changes.

# Specify Assured Recovery Properties

This section lists the Assured Recovery properties, corresponding values, and provides an explanation of each property.

**Note**: On Windows x64 systems, you cannot run scripts that activate applications with a graphical user interface.

**Scheduler**

The Scheduler enables you to automatically run Assured Recovery tests according to a pre-defined schedule, for example, every few hours, once a day, or several times a month. To set the scheduler, see Performing Assured Recovery Test in a Scheduled Mode.

**Start DB**

This property defines the first step in the AR test: starting the database services on the Replica.

**Automatic**

By default, this property is set to On. To use script to replace the automatic initiation of database services, set this option to Off.

**User-Defined Script**

You can specify a script to augment or replace the standard step to start the database services.

To replace the standard step, set **Automatic** to Off and set **User-Defined Script** to On. Then, specify the full pathname of the script to be executed in the **Script Name** box.

To execute the script following the standard step, leave **Automatic** set to On.

**Script Name (full path)**

Enter the name and full path of the script that is invoked following the starting of database services or instead of it.

**Arguments**

Additional arguments to pass to the script, which is specified in the previous property. Arguments entered here are static values.

**DB Testing of Replica**

This property defines the second step in the Assured Recovery test: verifying that all application services have started properly and that all databases or information stores have mounted successfully and are in a valid state.

**Automatic**

By default, this property is set to On. To use script to replace the automatic actions performed during this database validation stage, set this option to Off.

**User-Defined Script**

You can specify a script to augment or replace the actions performed during this database validation stage.

To replace the standard step, set Automatic to Off and set User-Defined Script to On. Then, specify the full pathname of the script to be executed in the Script Name box.

To execute the script following the standard step, leave Automatic set to On.

**Script Name (full path)** -- Enter the name and full path of the script that is invoked following the database validation step or instead of it.

**Arguments** -- Additional arguments to pass to the script, which is specified in the previous property. Arguments entered here are static values.

**Actions upon Successful Test (DB Online)**

After the Replica is successfully tested, the application data is in a known, valid state. You may want to make use of this fact, for example, to ensure that a backup is performed at this point on validated data. If the action you want to perform requires that the application is running and the databases or inform-ation stores are mounted, then it should be registered through a script here, in this step, by specifying the script details in the User-Defined Script boxes. This section has no default actions.

**User-Defined Script**

**Script Name (full path)** -- Enter the name and full path of the script that is invoked when the application is still running and the databases or inform-ation stores are mounted.

**Arguments** -- Additional arguments to pass to the script, which is specified in the previous property. Arguments entered here are static values.

**Stop DB**

This property defines the third and final step in a standard AR test: stopping the database services once the testing is complete.

**Automatic**

By default, this property is set to On. To use a script to replace the automatic stopping of database services, set this option to Off.

**User-Defined Script**

You can specify a script to augment or replace the standard step to stop the database services.

To replace the standard step, set **Automatic** to Off and set **User-Defined Script** to On. Then, specify the full pathname of the script to be executed in the **Script Name** box.

To execute the script following the standard step, leave **Automatic** set to On.

**Script Name (full path)** -- Enter the name and full path of the script that is invoked following the stopping of database services or instead of it.

**Arguments** -- Additional arguments to pass to the script, which is specified in the previous property. Arguments entered here are static values.

**Actions upon Successful Test (DB Off-line)**

As noted in Actions upon Successful Test (DB Online), the application is in a known valid state at this stage. You may want to copy it or perform a backup or take a snapshot at this time. If the action does not require the application to be running, register this through a script here, by specifying the full pathname of a script in the User-Defined Script field.

**Note**: On Windows Server 2003 and later systems, you can generate VSS snap-shots automatically. For more information, see Create VSS Snapshots Automatically.

**User-Defined Script**

**Script Name (full path)** -- Enter the name and full path of the script that is invoked after the Assured Recovery test is successfully completed.

**Arguments** -- Additional arguments to pass to the script specified in the Script Name property. Arguments entered here are static values.

## Assured Recovery Testing Limitations

When performing Assured Recovery Testing with Oracle or SQL Server databases, the software does not verify whether the database is actually mounted. It verifies only that the service is running. You can create a custom script that confirms both services are running and databases are mounted. Enable the appropriate user-defined script property. See Understanding Assured Recovery Properties for more information.

# Perform an Assured Recovery Test

The Assured Recovery test can be fully automated and performed on a scheduled basis as often as needed. Upon completion, appropriate personnel can be alerted with the status of the test, and additional actions can be triggered on success, for example, taking a VSS snapshot of the data or a running a backup. Alternatively, you can perform AR testing in non-scheduled mode, automatically or manually initiating the test whenever the need arises.

In both modes, the AR test is performed in steps, according to the AR configuration settings. Some of the steps are transparent, and are executed automatically whenever an AR test is performed. Other steps are visible, and can be configured as to whether and how they will be performed.

The standard steps are as follows:

1. Initiate Assured Recovery test - click the Replica Integrity Testing button on the toolbar to initiate the AR test on a scheduled basis or in a non-scheduled mode.

2. Suspend application of data changes on the tested Replica - this step is performed automatically at the beginning of each AR test.

3. Initiate a rewind component on the tested Replica - this step is performed automatically. It is aimed at capturing all the changes that are made to the Replica data during the test, so they can be later rewind back to the point when the replication was suspended.

4. Start the database services - by default, this step is performed automatically. However, it can be switched off, replaced, or followed by a user-defined script.

5. Test the database - the databases are verified, by default, using the same tests that are used to monitor the database in HA. These tests include verifying that all services have correctly started and that all databases have been successfully mounted. These tests can be switched off, replaced, or followed by a user-defined script.

6. Perform actions upon successful test while the database services are running - a user-defined script may be registered at this point to perform actions that are desired in the event of a successful test, but which also require that the application will be running.

7. Stop the database services - by default, this step is performed automatically. However, it can be switched off, replaced, or followed by a user-defined script.

Perform additional actions upon successful test while the database services are stopped - this step is optional, and it may be used to perform actions that take advantage of the fact that the application passed validation tests and that it was stopped in a systematic order.

8. Rewind AR Replica data and resuming replication - this step is performed automatically at the end of each AR test. It restores the Replica data to precisely the state it was before the test begun using the rewind technology. Then, it resumes replication.

# Performing Assured Recovery Test in a Scheduled Mode

When you set the AR test to run in a scheduled mode, it means that an AR test will be performed automatically on a regular basis. After you select this option the following flexible scheduling capabilities are offered:

- Testing on selected days of the week and for specific hours in a 24 hour cycle.

- Testing over selected periods (e.g., once every 36 hours) in a 7 day cycle.

- Exclusion of specific dates.

The AR test schedule can be set when the scenario is created or at a later stage.

**Note**: You can define only one scheduled task per scenario. If you attempt to configure AR testing while you already have a scheduled Suspend operation configured, the following message appears: **Only one scheduled task per scenario can be set. Suspend for host [Replica_name] is already switched on. Do you want to turn this option off now?** To switch the schedule option to the AR test, you need to click Yes.

**To schedule the AR test**

1. On the Scenario pane, select the Replica you want to test.

   On the Framework pane on the left, select the properties tab.

   The Replica Properties list appears.

2. If the scenario is running, click the Stop button on the Standard toolbar.

   The scenario is stopped.

3. On the Replica properties list, open the Scheduled Tasks group. Then, under the Replica Integrity Testing for Assured Recovery group, select the Scheduler property, and click the Not Set value.

   The Assured Recovery hours dialog appears:

   The Assured Recovery hours dialog is similar to the Schedule Setting dialog, which is used for scheduling automatic synchronization. For information about setting a schedule, see Schedule Synchronization.

4. Set the schedule for automatic AR testing in the Assured Recovery hours dialog, and click OK to save your schedule and close the dialog.

5. To activate the scheduler, click the Save button on the Standard toolbar and start the AR scenario.

The Replica you selected for testing will be tested on a regular basis according to the schedule you set.

# Performing Assured Recovery Test in a Non-Scheduled Mode

In a non-scheduled mode, you can test Assured Recovery either automatically or manually. When you are using the automatic method, all you need to do is initiate the AR test by a click of a button. Then, Arcserve RHA automatically performs all the test steps according to the AR configuration setting. Once the test is completed, the regular replication is resumed. There is only one difference between this method and a scheduled AR test. In a non-scheduled automatic mode, you initiate the test whenever you need, without using the Scheduler.

When you are using the manual method, you also need to initiate the AR test by a click of a button. However, unlike the automatic method, Arcserve RHA will suspend the test after the first standard step - starting the database service. This will occur even when all standard steps are configured as Automatic.

**Note**: If the Start DB property is set to Off, and there is no user-defined script that replaces it, the only thing that Arcserve RHA will do is suspend the application of changes to the Replica in preparation for the manual AR test.

Once the replication is suspended, you can perform tests or actions directly on the Replica without the need to later resynchronize the Master and Replica. You can use this option for manually testing applications or data on the Replica, or for performing tasks on the Replica instead of the Master, such as report generation, in order to reduce the Master workload.

When you finish the manual testing or operation, you need to manually stop the AR test suspension. This is done again by a click of a button. If other steps and actions were configured in the AR test, such as stopping the database services, they will be performed after you clicked the button for stopping the test and before the test will be declared as finished. When the test is considered finished, the replication will resume automatically.

# Perform Assured Recovery Test Automatically

**To perform Assured Recovery test automatically**

1. On the Manager, verify that the AR scenario is running.

2. To start the AR testing, on the Scenario pane select the Replica you want to test. Then, click the Replica Integrity Testing button on the Standard toolbar, or right-click the Replica and select Replica Integrity Testing from the short-cut menu.

   The Replica Integrity Testing for Assured Recovery dialog opens.

   In this dialog, the configuration you set for the AR test is displayed.

3. To start the automatic AR test using the existing configuration, click OK.

   **Notes**:

   - To change the test configuration before running the test, click Cancel, and refer to Setting Assured Recovery Properties.

   - To manually perform the AR test, select the Manual testing check box, click OK, and refer to Perform Assured Recovery Test Manually.

4. After you initiate the AR testing, the Replica Integrity Testing for Assured Recovery dialog is closed. Then, before the test begins to run, Arcserve RHA verifies that no synchronization, AR test or replication suspension tasks are currently in progress on any of the hosts that participate in the current scenario.

5. Once the verification stage ends, the AR test begins.

   The steps of the test are displayed as messages in the Event pane.

6. After the test is finished, the Replica is automatically restored to precisely the same state it was when the replication was suspended. This is done through the underlying rewind technology. Then, the changes that were accumulated in the spool are applied, and the replication is resumed and continues normally.

7. By default, once AR test is performed, an Assured Recovery Report is generated.

   **Notes**:

   - If the Assured Recovery Report was not generated, on the Replica Properties list under the Reports group, check the value of the Generate Assured Recovery Report property.

   - To view the report, refer to Viewing a Report.

All the tasks that were performed during the test are listed in the AR Report, along with their activation time and status.

# Perform Assured Recovery Test Manually

Instead of using the Scheduler, you can manually perform AR testing.

**To perform Assured Recovery test manually**

1. On the Manager, verify that the AR scenario is running.

2. To start the AR testing, on the Scenario pane select the Replica you want to test. Then, click the Replica Integrity Testing button on the Standard toolbar, or right-click the Replica and select Replica Integrity Testing from the pop-up menu.

   The Replica Integrity Testing for Assured Recovery dialog opens.

   In this dialog, the configuration you set for the AR test is displayed.

3. To start the manual AR test using the existing configuration, select the Manual testing check box. Once this check box is selected, the dialog changes to reflect only the actions that will be performed in a manual mode.

   **Notes**:

   ◆ To change the test configuration before running the test, click Cancel and refer to Setting Assured Recovery Properties.

   ◆ To automatically perform the AR test, clear the Manual testing check box, click OK and refer to Perform Assured Recovery Test Automatically.

4. Click OK to close the dialog and start the manual testing.

   ◆ If the Start DB property is set to On, or a user-defined script is set to replace it, these actions are performed and then the test is suspend.

   ◆ If no action is set to take place at this step, the replication and test are suspended at this stage.

5. From this stage, the only automatic action that Arcserve RHA performs, unless other actions are configured as Automatic, is suspension of updates on the Replica.

6. Once the replication is suspended, the following message appears in the Event pane: Replica is ready for Manual Integrity Testing.

Now, you can start performing any test you want directly on the Replica host, including making changes to the database. Note that these changes will not be saved once the AR test is finished, due to the rewind process.

**Important!** Do not restart the tested Replica at this stage. If you do, all the changes that accumulated on the spool will be lost.

7.  After you finished testing the Replica host, click again the Replica Integrity Testing button to resume replication.

    **Important!** If you do not click the Replica Integrity Testing button a second time at the end of the test, changes will continue to spool up on the Replica host. Eventually, the spool on the Replica host overflows and the scenario is stopped.

    A confirmation message opens.

8.  Click **Yes** to stop the AR test. If other steps and actions were configured in the AR test, such as stopping the database services, they will be performed before the test will be declared as finished. When the test is considered finished, the replication will be resumed automatically.

9.  After the test is finished, the Replica is automatically restored to precisely the same state it was when the replication was suspended. Then, the changes that were accumulated in the spool are applied, and the replication is resumed and continues normally.

10. By default, once AR test is performed, an Assured Recovery Report is generated.

# Chapter 13: Using VSS Snapshots

Arcserve RHA enables you to easily use Microsoft's Volume Shadow Copy Service (VSS) to create, view and manage VSS snapshots of the Replica data.

**Important!** You can use VSS only on Windows Server 2003 and up (not on earlier versions).

You can set up automatic creation of VSS snapshots in association with two operations: during replication suspension and after the Assured Recovery test is completed successfully. In addition, when Arcserve RHA is integrated with Arcserve, a VSS snapshot is automatically created upon each Arcserve Backup. All these snapshots are displayed in Arcserve RHA Snapshots Management window, which allows you to monitor and manage them.

This section contains the following topics:

# Create VSS Snapshots Automatically

By default, Arcserve RHA does not automatically create VSS snapshots. In order to activate this option, you need to set to On the **Create Shadow Copy (VSS)** property of the required Replica. This property is associated with two operations - replication suspension and AR test. Since you cannot set both operations on a scheduled mode for the same Replica, you need to configure the **Create Shadow Copy (VSS)** property with regards to one of these operations.

**Note**: Manual suspension cannot cause the creation of VSS snapshots. VSS snapshots will be created automatically only when associated with scheduled suspension.

# Setting Up Snapshot Creation

**To set up a snapshot creation**

1. On the Scenario pane, select the Replica for which you want to create VSS snapshots.

2. On the Framework pane on the left, select the Properties tab.

   The Replica Properties list opens.

3. If the scenario is running, click the **Stop** button on the toolbar. The scenario is stopped.

4. On the Replica Properties list, open the **Scheduled Tasks** group to display the **Suspend** and the **Replica Integrity testing for Assured Recovery** properties.

5. On either the **Suspend** or the **Replica Integrity testing for Assured Recovery** property, set the value to On.

   The **Create Shadow Copy (VSS)** property opens along with its related properties.

   You can set the VSS function switch in the Scheduled Tasks section.

   To change the properties of Create Shadow Copy (VSS), modify them from the Volume Snapshots Management Properties.

   **Notes**:

   ◆ If you set to On the **Replica Integrity testing for Assured Recovery** property, the **Create Shadow Copy (VSS)** property appears under the **Actions on Successful Test (DB Offline)** group.

   ◆ To associate VSS snapshot creation with the **Suspend** property, you need to schedule the suspension. Manual suspension will not create a VSS snapshot.

6. To activate the automatic creation of snapshots, set the **Create Shadow Copy (VSS)** property value to On.

7. Set the other VSS properties, according to the information provided in Understanding VSS Snapshot Properties.

8. Click the **Save** button on the Standard toolbar to save and apply your changes, and start the scenario.

   Now, after an AR test or during suspension, a VSS snapshot will be created automatically. The creation of the snapshot is indicated in the Event pane.

Once a snapshot is created, you can view and manage it through the Snapshots Management window.

# Understanding VSS Snapshot Properties

This section lists the VSS Snapshot properties, corresponding values, and provides an explanation of each property.

**Create Shadow Copy (VSS)**

To create VSS snapshots automatically during replication suspension or after successful AR test, set this option to On.

**Preferred Number of Snapshots to Keep**

Enter the number of snapshots you prefer to save and monitor. Once this number is reached, the oldest snapshots are replaced with newer ones. However, if the oldest snapshot is mounted or locked for backup, it is not deleted. Then, the new snapshot is added to the snapshot list even if the number is exceeded. Other internal VSS reasons can cause the number of saved snapshots to be higher than you specified. The default no. is 10 snapshots.

**Universal Shadow Storage Volume**

Specify the volume on which the snapshots will be stored. Note that this property cannot be set for each scenario separately. The storage location of the first VSS snapshot that is created in the system, applies to all other succeeding snapshots.

**Max Storage Size per Volume**

Enter the maximum storage allowed per volume used by snapshots (MB).

# Viewing and Managing Snapshots

Arcserve RHA provides you with a special window for managing your VSS snapshots.

This section contains the following topics:

- Viewing Snapshots
- Managing Snapshots

# Viewing Snapshots

**To open the Snapshots Management window**

- On the Manager, click the **Snapshot View** button on the Viewing toolbar.

  The **Snapshots Management** window opens.

  In this window, the VSS snapshots that were created for each existing Replica are displayed, according to the selected Replica.

  You can change the Replica whose snapshots are displayed by using the **Select replica host** drop-down list. The Replica hosts that appear on the list are all the Replica hosts that participate in existing scenarios.

  If a Replica that had snapshots participated in a scenario that was removed from the Manager, it does not appear on the list. To display snapshots of a Replica that no longer appear on the list, you can add it manually by using the **Add Host Name/IP** button.

  The following information is provided for each snapshot:

  - **Scenario name** - the scenario in which the snapshot was created.
  - **Snapshot Guid** - the unique ID that identifies the snapshot.
  - **Created** - the date and time of the snapshot creation.
  - **Creator** - the type of operation that is associated with the creation of the snapshot. Two types are available: Suspend and AR.
  - **Is Exposed** - indicates whether the snapshot was exposed ("True") or not ("False").
  - **Expose Path** - where the snapshot was exposed.
  - **Source Path** - the volume/directory that the snapshot captured.
  - **Storage Path** - where the snapshot was stored.
  - **Locked for Backup** - this column refers to snapshots that were taken as a part of Arcserve Backup. If the backup is not complete yet, you cannot manage the snapshot, and the value that appears is "True". If the backup is complete, or if the snapshot is not associated with Arcserve, the value is "False".

After the snapshots are displayed, you can start managing them.

# Managing Snapshots

**To manage snapshots**

- On the **Snapshots Management** window, select the snapshot you want to manage. Then, open the **Snapshot** menu and select the required option, or right-click and select the required option from the pop-up menu.

  The available actions are:

  - **Mount under Folder** - mount an exposed snapshot on an unused folder.

  - **Mount as Drive Letter** - mount an exposed snapshot on an unused drive letter.

  - **Unmount** - release an exposed snapshot without losing the snapshot itself. The snapshot is still exposed but it does not use a mount point.

  - **Delete** - delete a snapshot. You can delete several snapshots at once by using the **Ctrl** key.

  - **Refresh** - refresh the snapshot list to display the most up-to-date snapshots.

# Chapter 14: Using the Content Distribution Solution

This section provides instructions for creating, managing and using the Content Distribution solution.

This section contains the following topics:

# Understanding the Content Distribution Solution

**Important!** The Content Distribution solution requires a special license.

The Content Distribution solution is aimed at addressing the need of reliably spreading and managing information across a highly distributed environment. In a highly distributed IT environment, many servers contain the same or similar content that they receive from a single repository, and they simultaneously serve many end users. An example of such a distribute environment can be a large organization, which needs to deliver, synchronize and consolidate corporate information among internal users who reside in multiple locations and branch offices. This information can include price lists, policies, sales materials, manuals, and news. With the Content Distribution solution, employees and representatives in the field always have the right information at the right time.

The Content Distribution solution is also a powerful content delivery and web publishing solution that can serve your external customers. Through portals and web sites, you can deliver to your customers any information that is stored in files, from music to movies to documents to news. A good example is a service provider, who distributes content to dozens, hundreds or thousands of e-shops across the globe.

In a regular replication or HA scenario, the Master is usually the active or production server, while the Replica hosts are mainly a storage place for replicated data or standby servers. Unlike this role structure, in a CD scenario the Replica hosts are usually the active hosts, which directly provide information to end users, while the Master host only acts as the initial provider of updated data. The content is maintained in a single repository on the Master, and changes to the Replica hosts are delivered immediately or on a scheduled basis. When applying the CD solution to a large organization, multiple CD scenarios can use the same or overlapping root directories, apply different filtering options, and replicate data to different set of Replica hosts.

The CD solution is designed for one-to-many scenarios, meaning, scenarios that have one Master host and a large number of Replica hosts. These scenarios may replicate many files, or work with a small number of very large files. In this type of scenario, many of the Replica hosts are organized horizontally, as siblings on the same level, and not in hierarchal order as in parent-child relations.
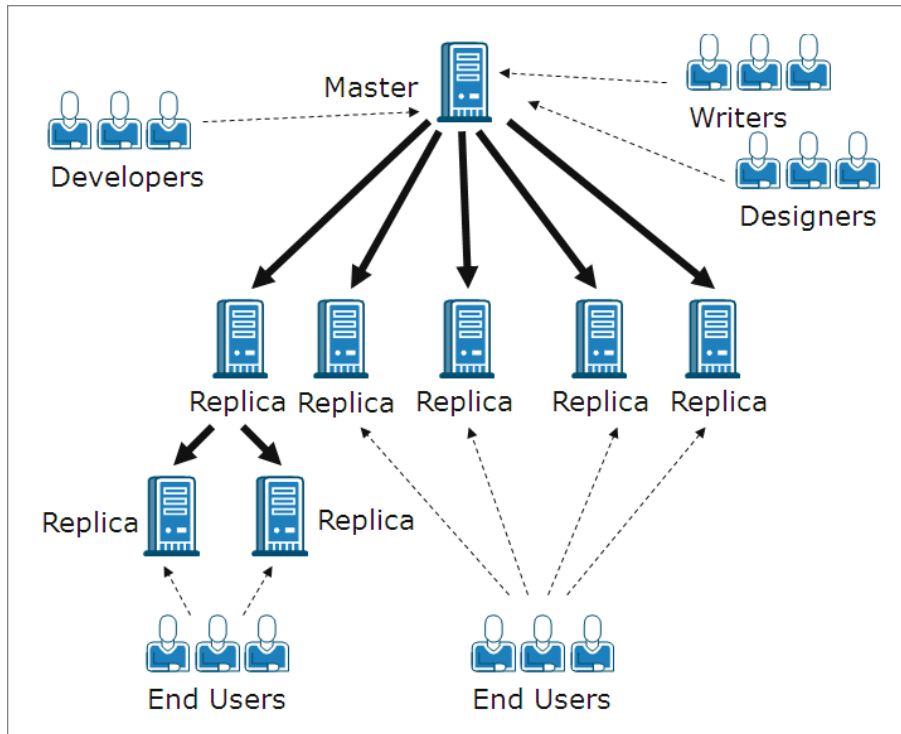
In a regular scenario that contains several Replica hosts on the same level, if more than one Replica host need re-synchronization following a reboot or some connection failure, all other Replica hosts will be re-synchronized as well. However, such a procedure might cause a performance problem when there are hundreds and thousands of Replica hosts. Therefore, in a CD scenario, if more than one Replica host needs re-synchronization, only the hosts that actually need synchronization will be re-synchronized.

Another feature of regular scenarios that might cause problems in a highly distributed environment, is the online replication mode. In a regular online replication mode, changes that occur on the Master are immediately transferred to the Replica, and overwrite the data that exist there. This process is useful for keeping the most up-to-date data on the Replica, but when users are directly using the data that is stored on the Replica, their work might be interrupted by the ongoing and continuous updates. To overcome this problem, a CD scenario can run in a special **On File Close** replication mode, a replication mode that is available only for CD scenarios.

In the **On File Close** mode, all data that is accumulated on the Master is transferred to the Replica, but it does not immediately overwrite the existing Replica data. Instead, data that was changed and transferred to the Replica is saved there as a temporary copy of the original file, and stored in a hidden directory. Once the original file is closed on the Master, the temp copy on the Replica is renamed. When the copy on the Replica receives the original file name, it replaces the older file

that is stored on the Replica, and keeps the data on the Replica updated. This method allows for an update process that does not interrupt the user's work. However, if the **On File Close** mode does not suit your environment needs, you can also use either the online or scheduled replication mode for your CD solution.

# Creating a Content Distribution Scenario

The creation of a CD scenario is similar to the creation of a Replication scenario for application and database servers. In both you are using the same step-by-step Scenario Creation wizard. The only major difference between them is that when you select the replication mode of a CD scenario, you have an additional replication mode. This replication mode, **On File Close**, is available only for Content Distribution purposes.

**Note**: This section demonstrates the configuration of a generic File Server Content Distribution scenario. For more detailed instructions involving scenarios tailored to specific applications, see the appropriate Operation Guide.

**To create a Content Distribution scenario**

1. Open the Arcserve RHA Manager. Then, select from the **Scenario** menu the

   **New** option, or click the **New** ![button icon] button on the Standard toolbar.

   The **Scenario Creation Wizard** opens.

2. Select the required scenario options, as follows:

   ◆ Select the **Create a New Scenario** option button.

   ◆ From the **Group** drop-down list, select the group to which you want to assign the new scenario, or enter a name for a new scenario group.

3. Click **Next**. The **Select Server and Product Type** page is displayed.

   A list of available applications and scenario types is presented.

   **Note:** The list of available applications depends on the licenses applied**.**

4. Select the required scenario options, as follows:

   ◆ From the **Select Server Type** list, select the type of server for which you want to create the scenario.

   ◆ From the **Select Product Type** options, select **Content Distribution Scenario.**

   **Note**: The **Tasks on Replica** options are not available for the CD solution.

5. Click **Next**. The **Master and Replica Hosts** page is displayed.

6. Enter the following information:

   ◆ In the **Scenario Name** box - accept the default name or enter a new name for the scenario. When entering a name, choose a unique name, since you cannot use the same name for more than one scenario.

> - In the **Master** and **Replica Hostname/IP** boxes - enter the hostname or IP address of the Master (source) and Replica (target) servers, or use the **Browse** buttons to find them.
> - In the **Port** boxes: accept the default port no. (25000) or enter new port numbers for the Master and Replica.

**Note**: If you want to include more than one Replica in the scenario, enter here the details of the first or most upstream Replica. After you finished the scenario creation, manually enter the other Replicas, as described in Add Additional Replica Servers.

7. [Optional] Select the **AssessmentMode** check box, if you want to gather statistics about the accurate bandwidth usage and compression ratio benchmarking without actually replicating data. If you select this option, no replication occurs, but a report is provided once the assessment process is completed.

8. [Optional] Select the **Verify Arcserve RHA Engine on Hosts** check box, if you want the system to verify whether Engines are installed and running on the Master and Replica hosts you specified in this page. If Engines are not installed on the selected hosts, you can use this option to remotely install the Engines on one or both hosts. For more information about the **Host Verification** page, see Creating a Scenario.

9. After you selected the desired options, click **Next**. The **Master Root Directories** page opens.

   Arcserve RHA displays the directories and files that are on the Master server. These directories and files are the data that can be replicated, protected and distributed. Arcserve RHA automatically aggregates data that has a common path into one directory.

10. Choose the directories and files you want to replicate from the Master to the Replica by selecting their check boxes. You can exclude folders and files from replication by clearing their check boxes.

    **Notes:**

    > - For more information about selecting and filtering root directories, refer to Creating a Replication Scenario.
    > - After you finish creating the scenario through the wizard, you can also select registry keys for synchronization, as described in Synchronize Registry Keys.

11. After defining the data to be replicated, click **Next**.

The **Replica Root Directories** page is displayed.

In this page you select the directories on the Replica where the replicated data will be stored.

**Important!** The Scenario Creation Wizard automatically configures the Replica root directories to be the same as the Master root directories. If you want to keep this configuration, ensure that your Replica server has the same drive letters as the Master server, and that the selected directories on the Replica do not contain data you want to save. You can change the default configuration at a later stage, as described on Select Replica Root Directories.

12. To change the Replica root directories, double-click the specified directories path. The **Browse and Select Replica Directory** dialog appears.

13. Select the directory on the Replica in which the replicated data will be stored, and click **OK**.

    You return to the **Replica Root Directories** page.

    **Note**: You can manually change the directory you selected for storing the replicated data, by clicking the selected directory name and entering a new directory. If you are entering a directory name that does not exist on the Replica, Arcserve RHA creates it automatically.

14. After defining the storage location of the replicated data, click **Next**.

    The **ScenarioProperties** page opens.

    The **Scenario Properties** page enables you to configure the scenario properties that affect the entire scenario. Typically, the default values are sufficient.

    If you want to configure the scenario properties at this stage, refer to Understanding Scenario Properties. To configure the scenario properties at a later stage, refer to Configuring Scenario Properties.

15. In the **Scenario Properties** page you can set the replication mode of the scenario. In addition to the two standard replication modes, **Online** and **Scheduling**, Arcserve RHA provides you with another replication mode that is especially designed for the CD scenario, the **On File Close** mode. The **On File Close** mode is similar to the **Online** replication mode with one difference: while in the **Online** mode data changes that are transferred from the Master to the Replica immediately overwrite the existing Replica data, in the **On File Close** mode changes to individual files will appear on the Replica only after the original file on the Master is closed. This way, if users are working directly with data that is stored on the Replica, their work will not be interrupted by constant updates.

**Note**: The default replication mode is **Online**.

To set the replication mode, open the **Replication** group and select the **Mode** property. Then, select the required replication mode from the drop-down list:

16. Once you set the scenario properties, click **Next**. The **Master and ReplicaProperties** page opens.

    The **Master and Replica Properties** page enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient.

    If you want to configure the Master and Replica properties at this stage, refer to Setting Master and Replica Properties. To configure the Master and Replica properties at a later stage, refer to Configuring Master or Replica Server Properties.

    **Note**: You can modify all the settings in this pane after the scenario is created. However, before changing any Spool properties (which can be configured here), review the Spool information for configuration details.

17. After you set the Master and Replica properties, click **Next**.

    Arcserve RHA verifies the validity of the new scenario and checks many different parameters between the Master and Replica servers to ensure a successful replication, distribution, and data recovery processes. Once the verification is completed, the **Scenario Verification** page opens.

    **Note:**Although Arcserve RHA allows you to continue with warnings, it is not recommended to do so. Resolve any warning situations before continuing to ensure proper operation of the application.

18. If the scenario is verified successfully, click **Next**.

    The **Scenario Run** page opens.

19. After the scenario is verified, you are prompted to run it. Running the scenario starts the data synchronization process.

    ◆ To add more Replica hosts to the scenario and to run it later, select **Finish**.

       **Note**: Arcserve RHA offers you two methods for adding Replica hosts to the scenario:

       — Through Arcserve RHA Manager, by manually adding each host to the scenario, as described in Add Additional Replica Servers.

– Through Arcserve RHA PowerShell, by using the **Add-Replica** and **Add-Replicas** commands. For more information about using Arcserve RHA PowerShell commands, refer to *Arcserve RHA Power-Shell Guide.*

◆ To run the scenario now, click **Run Now**.

The synchronization process starts.

20. Synchronization may take a while depending on the data size and network bandwidth between the Master and the Replica hosts. You will receive the following message in the Event pane once the synchronization is complete: **All modifications during synchronization period are replicated**. At this point, real-time replication is operational and the Content Distribution solution is installed and active.

**Note:** When the scenario has more than one Replica host, the **Scenario Statistics** tab does not display a graphic overview of the scenario state but scenario statistics organized in tables.

21. By default, once a synchronization occurs, a synchronization report is generated. For each Replica host that participates in the scenario, a separate Synchronization report is generated. For more information about opening reports, see Viewing a Report.

# Chapter 15: Managing Users

Arcserve RHA lets you manage a user's access rights through setting the content distribution scenario file ACL properties. ACL is Access Control List, a list of security protections that applies to the scenario file.

A special license is needed.

**Note**: You cannot manage ACL properties for replication or high availability scenarios.

This section contains the following topics:

# How Delegated Security Works

Delegated Security lets you control each user's access rights by setting the Content Distribution scenario file access control list (ACL) properties.

The ACL-based authentication model is a role based authentication model in Arcserve RHA. There are four pre-defined roles. Each role has pre-defined permissions defining what can be done to a scenario. The roles are:

- Super User
- Admin
- Control
- View-only

A Super User has full control rights for a scenario, while Admin, Control, and View-only have limited rights to the scenario. Only the Super User has the right to create a new scenario.

A user must be assigned one of the four roles to gain access to a scenario. A Super User or Admin can assign users or groups to any scenario and delegate rights to the users or groups. When a user tries to access a scenario through Arcserve RHA Manager or Arcserve RHA PowerShell, the current role is checked and, based on the permission of the role; the operation is allowed or denied.

# Access Rights Considerations

Before assigning user permissions, you should consider the following general conditions.

- All users are Windows domain or local users.

- A Super User has the right to create a new scenario.

- A Super User or Admin can assign users or groups of users to any scenario and delegate rights to the users or groups through Arcserve RHA Manager.

- The user or group list with the respective rights is stored in the standard NTFS ACL, applied to the scenario file.

- Super User can change the Super User group. However, after the change, the rights to all existing scenarios must be reassigned.

- Users are allowed to set or change a Super User group which is recorded in an encrypted registry. All Super Users belong to the group.

- The Access Control List is controlled indirectly via Control Service. Since Multiple GUI sessions may connect to one Control Service, the impersonation of each user account becomes indispensable.

# Prerequisite Tasks for Managing Users

You must perform the following prerequisite tasks before setting user rights or modifying user groups.

This section contains the following topics:

- Create a User Group
- Initial Group Selection
- Set Up a User Group

# Create a User Group

When using ACL authorizations, before you can open the Overview Page and the Manager, you need to create a Local Group. You need to define a Local Group with the name Arcserve RHA Users on the Control Service machine, and on all machines, that run the Arcserve RHA engine, where you want users or groups of users to have permission to add and modify replica hosts or access root directory of hosts.

**To create a Arcserve RHA User Local Group**

1.  On the Control Service machine, select Start, Settings, Control Panel, Administrative Tools, Computer Management.

    The Computer Management dialog opens.

2.  Select the **Local User and Groups** folder, and then open the **Groups** subfolder.

3.  Right-click on **Groups** and select **New Group**.

    The **New Group** dialog opens.

4.  In the **Group name** box, enter Arcserve RHA Users.

5.  To add the administrator user, click the **Add** button.

6.  Click the **Create** button to create the new Local Group, and then click the **Close** button to close the dialog.

    The new group is added to the Local Groups list on the Control Service machine.

**To use a customized name for User Local Group**

1.  Open the *mng_core_com.cfg* configuration file available in Control Service installation directory on the Control Service machine.

2.  Set the parameter *AclGroupName* value with the customized name.

3.  Restart the Control Service and then create user local group using the customized name.

## Initial Group Selection

You need to define a Local Group with the name Arcserve RHA Users on the Control Service machine, and on all machines that run the Arcserve RHA engine, where you want users or groups of users to have permission to add and modify replica hosts or access root directory of hosts.

When you open the Manager for the first time, the Manager detects whether a Super User Group already exists. If no Super User group is defined, the **Set Super-user Group** dialog appears.

The Set Superuser Group dialog displays the list of Groups that exist as local groups on the Control Service machine. You need to select the group that will include the members that will be defined as Super Users. You can later change this group.

# Set Up a User Group

To set up your environment to use ACL-base delegated security, Arcserve RHA lets you use your existing infrastructure or create a new network and local group. There are four groups required:

- Super User
- Admin
- Control
- View

You can assign users to a specific group depending on the required privileges for the individual user. For more information on user permissions, see Delegation of Rights.

**Note:** You can set up additional groups and designate them Super User, Admin, Control, View or use existing groups in the network.

On each computer participating in the scenario (Master, Replicas, and the Control Service), build a local group with the pre-defined name Arcserve RHA Users. Add groups and users of the organization to the Arcserve RHA Users local group as required.

When you open the user interface, if a Superuser group has not perviously been selected, you will be required to select one.

**Note:** Only a Super User can modify a Master server. Replica servers can be modified by a Super User, an Admin, or a Control role.

# How to Manage Users

Arcserve RHA lets you manage user permissions for scenarios by assigning individual users or groups delegated permissions.

The Super User or Admin manages users rights for a each individual scenario. From the user rights section of the user interface, you can set admin, control, or view permissions for a specific user or group for each scenario. This group or user then has the relevant permission for a particular scenario and can manage the scenario based on the rights they have been assigned. For example, a user or group can have user rights for one scenario and admin rights for an other scenario.

This section contains the following topics:

- Delegation of Rights
- Set User Rights
- Setting the Super User Group

# Delegation of Rights

User rights are set per user for using the Manager to make modifications to the Master host, or to the Replicas on its replication tree. User rights are assigned per scenario.

You can assign user permissions based on the following:

| Operation | Super User | Admin | Control | View only |
|---|---|---|---|---|
| Set user rights | Yes | Yes | No | No |
| Edit Master host | Yes | No | No | No |
| Edit replication mode | Yes | Yes | No | No |
| Edit schedule mode | Yes | Yes | Yes | No |
| Edit Master spool size | Yes | No | No | No |
| Modify reports on Master | Yes | Yes | No | No |
| Edit replica host | Yes | Yes | Yes | No |
| Edit replica spool size | Yes | Yes | No | No |
| Run a scenario | Yes | Yes | Yes | No |
| Stop a scenario | Yes | Yes | Yes | No |
| Synchronize a scenario | Yes | Yes | Yes | No |
| Restore data | Yes | Yes | No | No |
| Modify Master notification | Yes | Yes | No | No |
| Modify Replica notification | Yes | Yes | No | No |
| Generate a difference report | Yes | Yes | Yes | Yes |
| Set a bookmark | Yes | Yes | Yes | No |
| Show difference report | Yes | Yes | Yes | Yes |
| Run high availably resources | Yes | No | No | No |
| Check a scenario state | Yes | Yes | Yes | Yes |
| Suspend a replica | Yes | Yes | Yes | No |
| Modify reports on replica | Yes | Yes | No | No |
| Modify Master trigger file | Yes | Yes | Yes | No |
| Modify Replica trigger file | Yes | Yes | Yes | No |

# Set User Rights

You can set or reset user rights for a specific scenario.

**Note:** For all scenarios with licenses other than a delegated security license, you must reset the user rights.

**To set user rights**

1. From the Arcserve RHA> Manager Scenario menu, select Scenario, User Rights.

   **Important!** The Scenario menu contains the extra option User Rights. This option is only available to users with Super User or Admin rights.

   The Security window opens displaying the security rights of each scenario.

2. Click Add.

   The Select Users or Groups window opens.

3. From the Look in field drop-down list, select a Domain.

4. Select the required user or group.

   **Note:** Multiple user or group selections are not supported.

5. Click Add and then click OK.

6. From the Permission column, set access rights for a user or group from the drop-down list.

   **Important!** If you remove yourself (Admin) from the list in the security window, you will no longer have any user rights in the current scenario. After restarting the Arcserve RHA Manager or waiting for approximately 10 seconds, the scenario no longer appears in the list of scenarios.

# Setting the Super User Group

You can change a Super User group at any time.

**To modify the Super User group**

1. On the Manager, open the **Scenario** menu and select the **Set Superuser Group** option.

   The **Set Superuser Group** opens.

2. From the **Groups on Control Service** list, select the group to which you want to assign the Super User group.

# Chapter 16: Managing Services

Arcserve RHA also provides a mechanism for automating the management and monitoring of the services critical for application availability. Services management is built into the Scenario Creation Wizard and can also be manually accessed from the Arcserve RHA Manager Root Directories tab.

The Services Management capability is designed to provide a framework for protecting applications that cannot be protected by dedicated Arcserve RHA scenarios (for example, Microsoft SQL or SharePoint Server). Rather than write custom scripts to manage services, Arcserve RHA can start, stop, and trigger switchover based on the status of the services you specify.

**Note:** This feature is not applicable for file server scenarios.

This section contains the following topics:

# 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**
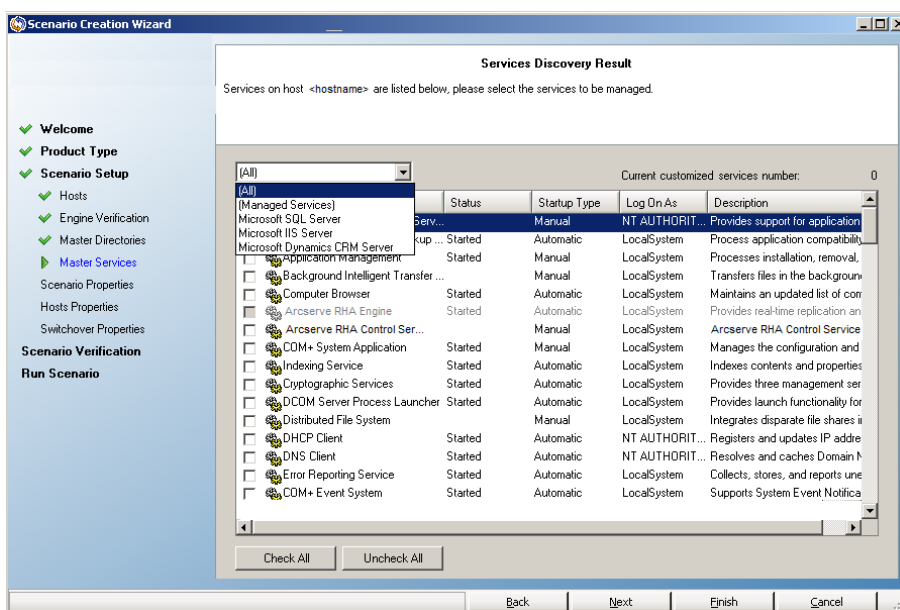


— **All** - lists all services discovered on the Master server

— **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.



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 17: Managing Clusters

Arcserve RHA now supports non-shared disk clusters, extending native high availability inherent in cluster environments from LAN to WAN.

This section contains the following topics:

# Introduction to Windows 2008 Cluster Enhancements

Effective with the release of Microsoft Cluster Services (MSCS) 2008, the cluster itself now supports nodes in different physical locations with or without shared storage and manages its own failover. However, the shared disk that stores cluster data can be a single point of failure should it suffer damage or data loss.

Arcserve RHA provides data replication for Shared Storage Clusters.

# How Arcserve RHA Enhanced Cluster Support Works

With the new Arcserve RHA plug-in for Microsoft Failover Cluster (MSFC) 2008, you can copy data between cluster nodes. There are three node roles:

- **Source Node** -- This is the node in the cluster currently communicating with the disk resource you wish to protect.

- **Destination Node** -- This is the node where you wish to copy data.

- **Inactive Node** -- This is the node in a cluster that is not communicating with the disk resource at a given point in time.

To copy data between cluster nodes, first create a new resource, an RHA disk to store data replicated from the source node. You can manage the new disk resource with the MMC snap-in, if you wish to change properties or monitor replication status. The Arcserve RHA Control Service (which should be installed on a separate server) is used to apply licenses. MSFC manages failover within the cluster while Arcserve RHA replicates the disk resource.
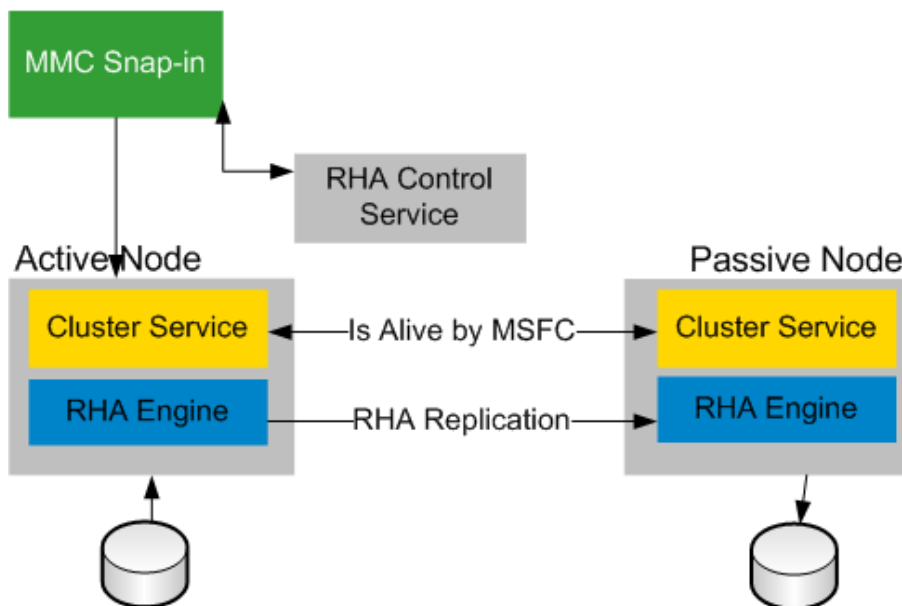
# Deploy Arcserve RHA Components for Cluster Support

This section contains the following topics:

- [Install Arcserve RHA for Microsoft Failover Clusters](#)
- [Launch Arcserve RHA for Microsoft Failover Clusters](#)
- [Console actions](#)

# Install Arcserve RHA for Microsoft Failover Clusters

To install Arcserve RHA for Microsoft Failover Clusters, refer to the Arcserve RHA Installation Guide.

# Launch Arcserve RHA for Microsoft Failover Clusters

After installation, you may launch the software from the Windows Start menu.

You may also manually integrate the software with the Microsoft Failover Cluster Manager.

**To manually integrate Arcserve RHA for Microsoft Failover Clusters**

1.  Click Start, Run and launch mmc.exe.

2.  From the Console menu, click File, Add/Remove Snap-in.

3.  From the list of available snap-ins, choose the Arcserve RHA Console for MSFC and Add it to the Selected Snap-ins list.

4.  Choose the Failover Cluster Manager and Add it to the Selected Snap-ins list.

5.  Click OK.

The Console Root tree updates to include the specified snap-ins. From here, you may manage your Arcserve RHA disk resources.
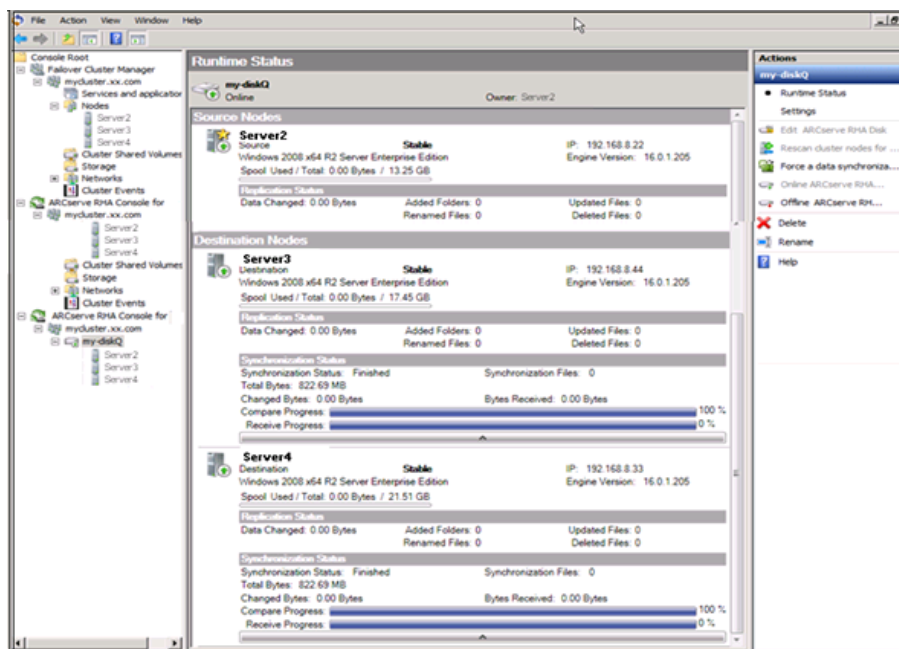
# Console Actions

From the Arcserve RHA for Microsoft Failover Cluster Console, you can perform various actions to help you manage and monitor replication status. The Actions available depend on the selection you make from the Console. You may also access these options by right-clicking a level in the Console structure.

**Console Level**

- **Open Cluster** -- Provide a cluster name and click OK.

- **View**-- Click Customize to select the components to display on the Console screen and click OK.

- **New Window from Here** -- Opens a new Console window.

- **Refresh**-- Updates the display.

- **Help**-- Opens Console Help.

**Cluster Level**

- **Add Arcserve RHA Disk** -- Opens the Add Disk dialog. Provide a Disk Name and optional volume and click OK.

- **View** -- Click Customize to select the components to display on the Console screen and click OK.

- **New Window from Here** -- Opens a new Console window.

- **Refresh** -- Updates the display.

- **Help** - Opens Console Help.

**Disk Resource Level**

- **Runtime Status** -- The Console lets you check Runtime Status of the disk resources in your cluster. Information including Replication Status, Engine Version, and Synchronization Status is provided.

- **Settings**-- You may set the following properties:

  - Synchronization Type -- File or Block Synchronization. File Synchronization is best suited for File Server applications and Block Synchronization is preferred for database applications.

  - Ignore Files of Same Size/Time -- On or Off. This setting is a filter that allows data comparison to consider only file size and modification time when determining whether two files differ. When set to On, files with the same path, name, size, and time are skipped. You should disable this property when Synchronization Type is Block Sync.

- Replicate NTFS Compress Attribute -- On or Off. This setting replicates the compress attribute in files or directories during synchronization and replication.

- Replicate ACL -- On or Off. This setting replicates ACLs for files and directories during synchronization or replication.

- Replicate NTFS ADS -- On or Off.

- Number of Streams -- Defaults to 1. In WAN environments with high bandwidth and latency, you may increase the Number of Streams to improve utilization. Click the drop-down arrow to set the value or specify it in the ws_rep.cfg file.

- **Edit Arcserve RHA Disk Resource** -- Rename an existing disk resource, set the optional volume or change it and click OK. You must take the disk resource offline to edit it.

- **Rescan cluster nodes for this resource**-- Do this when you add or eject host nodes to and from the cluster.

- **Force a data synchronization** -- Lets you perform immediate synchronization.

- **Online Arcserve RHA Disk Resource** -- Lets you bring a disk resource online. Synchronization between source and destination nodes is automatically performed and data changes on the source node are replicated in real time to the destination node. Progress is displayed on the Runtime Status pane.

- **Offline Arcserve RHADisk Resource** -- Lets you take a disk resource offline so you can edit it. Data replication is stopped while the dis resource is offline.

- **View** -- Group by Role or Group by Topology. Role lets you view disks by status. Topology lets you see which disks are assigned to which sites.

- **New Window from Here** -- Opens a new screen.

- **Delete** -- Lets you remove a disk resource.

- **Rename** -- Lets you change the name of the disk resource. The highlight bar changes to a rectangle and a cursor is enabled.

- **Help** -- Opens Cluster help.

**Node Level**

- **View** -- Click Customize to select the components to display on the Console screen and click OK.

- **New Window from Here** -- Opens a new Console window.

- **Help** -- Opens Console Help.

# Open a Cluster

**To open a cluster, you must select the Arcserve RHA for MFC entry in the Console Root list.**

1. Click Arcserve RHA Console for MFC in the Console Root list.

2. From the Actions list, click Open Cluster.

3. On the Open Cluster dialog, enter the cluster name and click OK.

    The cluster is listed in the License box.

# Customize View

The View Action lets you select the components to display on the console screen. The same View Action is available from all Console Root levels.

**To customize your view**

1. From the Console Root list, select the Console, Cluster, Disk Resource or Server to highlight it.

2. From the Actions list, click View.

3. Click Customize.

4. Click to select or clear the MMC and Snap-in components you wish to show or hide and click OK.

Your view is changed accordingly.

# Open a New Window

You can open additional console windows as needed. This action is available from all Console Root levels and functions like the Microsoft command.

**To open a new window**

1. From the Console Root list, click the console, cluster, disk resource or server.

2. From the Actions list, click New Window from Here.

 A new window opens. Use the standard Windows controls to minimize, maximize and move it as needed.

# Refresh the Display

You can update the display from all console selection.

**To update the display**

1. Click Refresh from any console level.

# Get Help

You can obtain console Help from any console selection.

**To get help**

1.  Click Help from any console selection.

## Add Arcserve RHA Disk Resource

You can create a new Arcserve RHA disk resource only from the cluster level in the Console Root list. See Create New Disk Resource.

# See Runtime Status

You can view the status of the disk resources in your cluster and obtain information such as replication status, engine version, and synchronization status.

Runtime Status is an action available only when the disk resource is selected from the Console Root list.

**To see Runtime Status**

1. Click the disk resource in the Console Root list.

2. Click Runtime status.

   The status is displayed by source and destinations roles, by default.

3. (Optional) Click View, Group by Topology to see status grouped by Sites.

**Note**: The Group by Role and Group by Topology selections are only available from the Disk Resource View action.

# Configure Server Settings

You can set the following properties.

- **Synchronization Type** -- File or Block Synchronization. File Synchronization is best suited for File Server applications and Block Synchronization is preferred for database applications.

- **Ignore Files of Same Size/Time** -- On or Off. This setting is a filter that allows data comparison to consider only file size and modification time when determining whether two files differ. When set to On, files with the same path, name, size, and time are skipped. You should disable this property when Synchronization Type is Block Sync.

- **Replicate NTFS Compress Attribute** -- On or Off. This setting replicates the compress attribute in files or directories during synchronization and replication.

- **Replicate ACL** -- On or Off. This setting replicates ACLs for files and directories during synchronization or replication.

- **Replicate NTFS ADS** -- On or Off.

- **Number of Streams** -- Defaults to 1. In WAN environments with high bandwidth and latency, you may increase the Number of Streams to improve utilization. Click the drop-down arrow to set the value or specify it in the ws_rep.cfg file.

# Edit Arcserve RHA Disk Resource

You may rename an existing disk resource or set and change the volume.

**Note**: You must take the disk resource off-line to edit it.

**To edit a Arcserve RHA disk resource**

1. From the Console Root list, select the disk resource you wish to edit.

2. Click Offline Arcserve RHA disk resource.

    Wait while the operation is completed.

3. When the Edit Arcserve RHA Disk Resource Action changes from gray to black, click it.

4. Enter a new name or volume, as desired, and click OK.

5. Click Online Arcserve RHA Disk Resource to bring the disk back online.

# Rescan Cluster Nodes

You must reconfigure your Arcserve RHA disk resource if you have added or ejected host nodes to and from the cluster.

**To rescan cluster nodes**

1. Select the disk resource from the Console Root list.

2. Click Offline Arcserve RHA Disk Resource to halt replication.

3. Click Rescan cluster nodes for this resource.

4. Click Online Arcserve RHA Disk Resource to restart replication.

# Force Synchronization

Synchronization is performed automatically whenever you bring a disk resource online. However, you can perform synchronization at will, such as prior to a hardware replacement.

**To force synchronization**

1. Click the disk resource in the Cluster Root list.

2. From the Action list, click Force a data synchronization.

# Bring a Disk Resource On or Off-Line

After you create a disk resource, you must apply the license and then bring it online, which synchronizes the entire source node with the destination. After Synchronization is completed, all data changes are replicated in real time to the destination node. Replication status is readily available in the Runtime Status view.

To change a disk resource, you must take it off-line. Replication is temporarily suspended.

**To bring a disk resource on and off-line**

1. Select the disk resource you created from the Console Root list.

2. Click Online Arcserve RHA Disk Resource.

3. Click Runtime Status. Wait while status changes to Online and Stable. Synchronization Status starts at Initializing and when finished, changes to Finished.

4. If you wish to bring the disk resource off-line, click Offline Arcserve RHA Disk Resource.

5. Wait while Replication is halted.

6. Click Runtime Status, verify that status is Offline and Stopped.

   The disk resource is off-line.

   When you are ready to return the disk resource to its former status, select it and click Online Arcserve RHA Disk Resource. If desired, you may also Force synchronization.

# Replicate Cluster Data and Manage Resources

Cluster storage can be a single point of failure. To protect cluster data, use Arcserve RHA for Microsoft Failover Cluster to replicate data to another disk. This disk can be remote.

**Viewing Clusters**

- Open a cluster
- Browse Cluster
- Customize the View
- Open a New Window
- Refresh the Display

**Protect Cluster Storage**

- Add Arcserve RHA Disk Resource
- Edit Arcserve RHA Disk Resource
- Rescan Cluster Nodes
- Force Synchronization
- Configure Server Settings
- See Runtime Status
- Bring a Disk Resource On- or Off-Line

**Manage Licenses**

# Open a Cluster

After your cluster is configured you can manage it using the Arcserve RHA for Microsoft Failover Cluster console.

**To open a cluster**

1. From the console, right-click the Arcserve RHA Console for MSFC.

2. Choose Open Cluster.

3. Specify the Cluster Name and click OK.

   The cluster you specified is loaded and displayed in the panel.

**Note**: You can also open clusters from the Actions list on the right side of the console.

# Browse Cluster

If you wish to see what disk resources are contained in a specific cluster, or what cluster nodes use a specific disk resource, you can browse clusters. You may also determine the status of the disk resource and the cluster host.

**To browse a cluster**

1. From the Microsoft Management Console, expand the Root Scope Node in the scope node panel. All disk resources are listed under the cluster node.

2. Expand the Arcserve RHA Disk Node. All cluster host nodes are displayed under the disk resource.

# Create new Disk Resource

Using theArcserve RHA MMC snap-in, you can create a new Disk Resource which acts as the Replica in a Arcserve RHA scenario. The Arcserve RHA Disk Resource is not supported for use a quorum disk.

Whenever you change a cluster node (eject, add node), you should rescan cluster nodes for that resource. For more information, see the topic, Rescan Cluster Nodes.

If the real disks are shared by several cluster nodes, make sure the same drive letter is used on each cluster node when you create the disk resource. If the real disk is offline, bring it online from the Microsoft Disk Management Console, check whether the expected drive letter was assigned to the volumes that belong to the disk. You can change the drive letter from the Microsoft Disk Management Console.

**Note**: You may create the disk resource without completing the Target Volume field, but must specify this information to bring the disk resource online. If you wish to add the volume information later, click the option, Edit Arcserve RHA disk resource.

**To create a new disk resource**

1. From the Arcserve RHA for MSFC console, click the cluster node.

2. In the Actions list, click Add Arcserve RHA Disk Resource.

3. Specify the Disk Name and Target Volume and click OK.

The disk resource is created under the cluster node in the console. Nodes using the disk resource are automatically detected and listed.

**Next Steps:**

- Manage Licenses

- Add Arcserve RHA Disk Resource to Service or Application

# Add Disk Resource to Applications

After you create a Arcserve RHA Disk Resource in the Arcserve RHA for MSFC Console, create a new application with the Configure an Empty Service or Application Wizard in Microsoft Failover Cluster Manager and add the disk resource you created.

# Manage Licenses

Before you apply a license, you should first create the Arcserve RHA disk resource. License count should match the number of nodes in the cluster, not the number of Arcserve RHA disk resources. You do not need to apply the license every time you create a disk resource.

**To manage licenses**

1. From the Arcserve RHA for MFC console, click Edit License Server to assign a license server. If the license server has already been assigned, skip to Step 5.

2. Provide the IP Address and Port Number of the license server, which is the server on which the Arcserve RHA Control Service is installed.

3. Click OK.

4. Provide the User Name and Password and click OK.

5. Click Update License. The Update License dialog opens. Provide the user name and password and click OK.

The License Count is updated.

**Note**: License reservations must be manually cancelled before you delete a cluster or the license key could be lost. Click Update License to release existing licenses.

# Browse Events

In the MMC interface, choose Cluster Events to view the logs.

On the Cluster Events Filter dialog, you can filter events based on the following:

- Nodes
- Event Logs
- Level
- Event ID
- Date range

# Chapter 18: Creating and Executing User-Defined Scripts

Though scenarios are flexible, easy to use and intuitive to create, there may be times when you need powerful customization options to extend software benefits. User-defined scripts provide such customization, allowing you to embed additional operations in your scenarios. Scripts are limited only by your operating system. If a file can be executed from the command line on a particular host, it can be applied as a script within a scenario.

Scripts work in all product releases.

You can use batch files (.bat or .cmd), VBScript (.vbs) with Cscript.exe, or Power-Shell (.ps) scripts with Arcserve RHA. Arcserve RHA r12.x and later versions include a PowerShell snap-in. Shell scripts (.sh) can be used on UNIX systems.

## Example

You can create a batch file (example.bat) that executes a VBScript file with Cscript. To do so, first call the Cscript executable and then pass the VBScript file as a file call. Specify the batch file name in the Script Name (full path) field in the appropriate scenario property.

This section contains the following topics:

# How User-Defined Scripts Work with Arcserve RHA

Scripts for use within Arcserve RHA and its prior releases must be written to return numeric codes that define success (0) and failure states (any non-zero value). The software displays return codes in the event window so you can determine immediately when and where failures occur.

Scripts follow this format:

ScriptName (full path) Arguments

The script name is the name and full path of the executable script to be invoked. Add directories to this property in the form of <drive>:\<dir>\<file.ext>. The software displays directories as <drive>:/<dir>/<file.ext>. Arguments passed to the script are static and literal values.

Scripts must have the same name and reside in the same directory on both the Master and Replica servers.

**Note**: You cannot run scripts that activate user interface applications on Windows x64 systems.

Scripts can be executed from different scenario properties. The following topics describe the properties and how to set scripts for each.

# User-Defined Script Properties

You can execute user-defined scripts within the following properties panels:
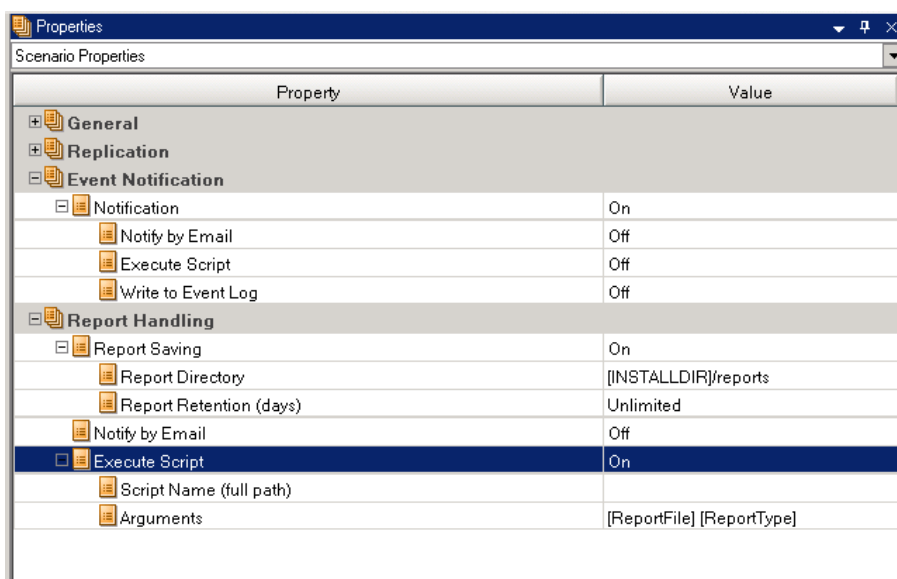
- **Scenario Properties** -- Event Notification, Report Handling

- **Master Properties** -- Replication, Event Notification, Report Handling

- **Replica and Scheduled Task Properties** -- Replication, Scheduled Tasks, Event Notification, Report Handling

- **High Availability Properties** -- Network Traffic Redirection, Is Alive, DB Application, Action Upon Success

- **Assured Recovery Properties** -- Action Upon Successful testing (DB Online, DB Offline)

# Execute User-Defined Scripts From Scenario Properties

You must save scripts to be executed from Scenario Properties on the Control Service host.

**Note**: Scripts added to Scenario Properties affect the scenario, the Master, and the Replica hosts. If you add a Reports Handling script here as well as on one of the hosts in the scenario, duplicates occur.

You may execute scripts from the following property groups:



- **Event Notification** -- this script permits you to handle events or errors as they occur. Turn On the Execute Script property. Enter the name in the Script Name (full path) field. Provide arguments to be passed to the script in the Arguments field.

- **Report Handling** -- this script is invoked after a report is generated. Turn on the Execute Script field. Enter the name in the Script Name (full path) field. Provide arguments to be passed to the script in the Arguments field.

To specify a Scenario Properties script, see the topic, Specify a Custom Script in a Property.

# Execute User-Defined Scripts from Master Properties

Master properties let you specify scripts that act on the Master server.

You may execute scripts from the following property groups:



**Replication**

— If you turn On the Run Script Before Synchronization property , the synchronization process does not start until this script completes.

— If you turn On the Run Script After Synchronization property, the script runs on the Master immediately after synchronization begins. Synchronization does not wait for the script to complete.

— If you turn On Run Script upon Trigger File Creation property, (File Server scenarios only), the special actions defined in the script execute when the specified trigger file appears.

**Event Notification** -- this script permits you to handle events or errors as they occur. Turn On the Execute Script property. Enter the name in the Script Name (full path) field. Provide arguments to be passed to the script in the Arguments field.

**Report Handling** -- this script is invoked after a report is generated. Turn On the Execute Script property. Enter the name in the Script Name (full path) field. Provide arguments to be passed to the script in the Arguments field.

To specify a script in Master Properties, see the topic, Specify a User-Defined Script in Properties.

# Execute User-Defined Scripts from Replica Properties

You may execute user-defined scripts that run on the Replica server from the following property groups:



**Replication**

— **Run Script Before Synchronization** -- turn On this property to execute a script that runs on the Replica immediately before synchronization. Synchronization does not start until the script completes and can be used for starting certain third-party services.

- **Run Script After Synchronization** -- turn On this property to execute a script that runs on the Replica immediately after synchronization begins. It does not wait for synchronization to complete.

◆ **Event Notification --** this script provides a way to customize the handling of events and errors. Turn On the Execute Script property. Enter the name in the Script Name (full path) field. Provide arguments to be passed to the script in the Arguments field.

**Report Handling --** this script is executed after a report is generated. Turn On the Execute Script property. Enter the name in the Script Name (full path) field. Provide arguments to be passed to the script in the Arguments field.

To specify a script in Replica and Scheduled Task Properties, see the topic, Specify a User-Defined Script in Properties

# Execute User-Defined Scripts from Scheduled Task Properties

You may execute user-defined scripts that execute upon successful testing:



**Scheduled Tasks -- Replica Integrity Testing for Assured Recovery**

 — **Start DB** -- If the Automatic property is set to On, the Start DB property determines the first step in the AR process, starting database services on the Replica. The script runs when services are started and the database is

mounted. AR does not continue until the script completes. If the Automatic property is set to Off and User-Defined Script to On, you may then specify a script to replace the standard first step.

— **DB Testing of Replica** -- If the Automatic property is set to On, the DB Testing of Replica property determines the second step in the AR process, verifying that all application services started properly and that all databases are mounted successfully and in a valid state. For example, this property could be used to execute a DBCC Check on a SQL Server to verify the data set. Integrity testing does not continue until the script completes and if it fails, the AR test fails, as well. If the Automatic property is set to Off, and the User-Defined Script to On, you can replace this standard second step.

— **Action Upon Successful Testing (DB Online)** -- After the Replica is successfully tested, the data is in a known valid state. This property lets you make use of this knowledge. For example, you could ensure that a backup is performed at this point, ensuring it is done on validated data. Or, you could ensure that an application is running.

— **Stop DB** -- This property determines the final step in the standard AR test, stopping the database services after the test is completed. To replace the standard test, turn Off Automatic and turn On User-Defined Script. You should use this property to stop anything started by a script in the Start DB property.

— **Action Upon Success (DB Offline)** -- After the Replica is successfully tested, the data is in a known valid state and you may wish to copy it, perform a backup, or take a snapshot. If the action you wish to perform does not require the database to be running, use this property to register your script. If you are running Windows Server 2003 (or later), you can generate VSS snapshots automatically.

To specify a script in, see the topic, Specify a User-Defined Script in Properties

# Execute User-Defined Scripts from High Availability Properties

You can execute scripts to run in HA scenarios from the following property groups:
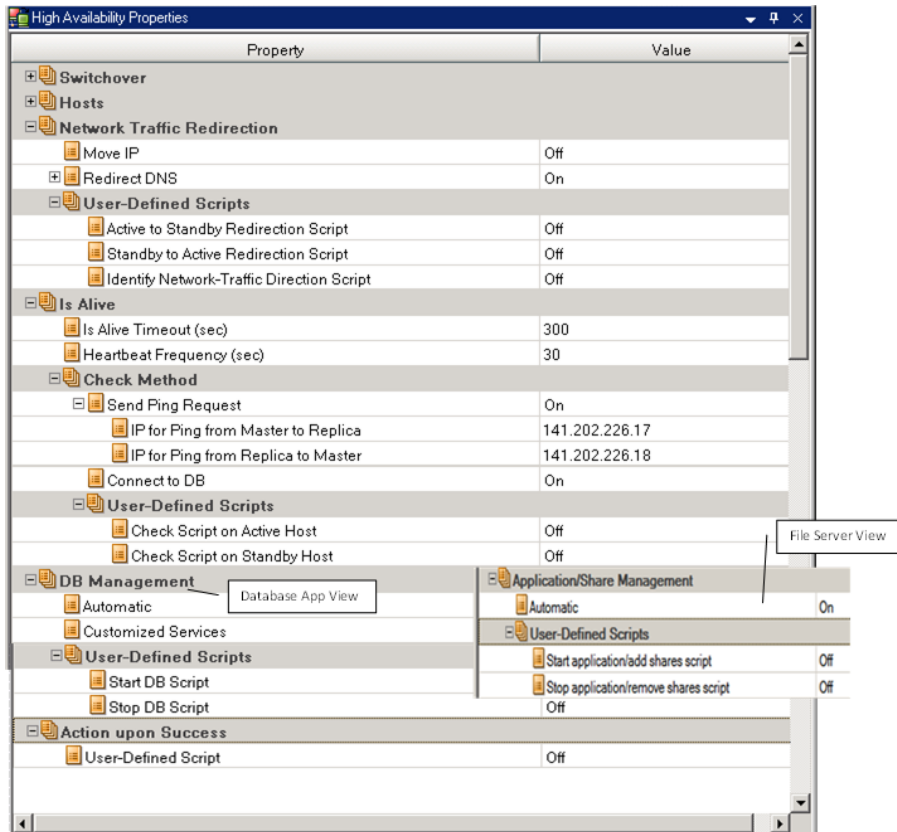


- **Network Traffic Redirection**

    — **Active to Standby Redirection** -- Turn On this property to execute a script that runs on the active server, if it is available, to redirect users to the standby host, or to release network resources on the active host.

    — **Standby to Active Redirection** -- Turn On this property to execute a script that runs on the standby server, to redirect users to the active server.

    — **Identify Network Traffic Redirection** -- Turn On this property to execute a script that determines which server is active. If the script returns 0, the host it was executed from is considered to be the active server. If the script returns a non-zero value, that host is considered inactive.

- **Is Alive, Check Method**

  - **Check Script on Active Host** -- this script runs on the active server during every Is Alive heartbeat to verify that the server is available to users.

  - **Check Script on Standby Host** -- this script runs on the standby server during every Is Alive heartbeat to verify that the server is available to users.

  - **Start DB/Start Application/Add Shares** -- When set to On, the script augments or replaces the start of DB services and applications, or enables folder sharing. The action defined occurs during the scenario run process on the active host, or during a switchover on the standby host.

  - **Stop DB/Stop Application/Remove Shares** -- When set to On, the script augments or replaces the stopping of DB services and applications, or disables folder sharing. The action defined occurs during the scenario run process on the standby host, or during a switchover on the active host.

- **Action Upon Success** -- When set to On, the script executes its defined actions after a successful switchover.

# Specify a User-Defined Script in Properties

The following procedure explains how to enable user-defined scripts from the various properties. You are not limited to one script. You may specify scripts for various purposes but exercise caution. Report Handling and Event Notification scripts can be specified in several property groups, which could result in duplicate execution.

**To specify a custom script in a property**

1. Stop the scenario if it is running.

2. Select the appropriate Properties tab for the desired server. For more information, refer to the topic, User-Defined Scripts Properties.

3. Expand the desired Property Group.

4. Set the appropriate property according to the instructions provided in the Execute User-Defined Scripts topics.

5. Resume running of the scenario.

# Troubleshoot Script Use

**I receive an error code (1) when my VBS script using cscript.exe is executed.**

Error code (1) from a script means that there is an invalid parameter or that the script cannot be found. Check the arguments in the scenario and verify that the syntax is correct and that all characters, especially the quote symbols, are present. Sometimes, when copying and pasting into the arguments field within the scenario, the quotes can be taken as literal special characters and are not presented to cscript.exe correctly at execution.

**I receive an error message, ER00160: Script<script name> did not finish execution in <number> "seconds" during switchover.**

This problem can occur when a high availability script takes longer than the designated timeout period to complete. The default time is 300 seconds (5 mins). You can modify the ws_rep.cfg file within the Engine install directory to adjust this time. The attribute to modify is HAScriptExecutionTimeout=300. To change this value, remove the # in front of the attribute, modify the right-hand numeric value, and then restart the engine service.

**I receive the error, ER00564: Script<script> doesn't exist. Or, ER00569: Script<script> not found, is displayed.**

This means that the script that was supposed to execute is not in the location specified. Check that there are no spaces within the root directory. If there are spaces, the script directory needs to be placed inside quotes, like "c:\Program Files\ScriptExamples".

**I receive the error, Check script on standby host failure, or Check script on Active host failure.**

This means that the Is Alive script returned a fail code on the host specified and a switchover will either take place or needs to take place.

**When starting the scenario, I receive an error, ERROR: Network traffic is directed to host<host>, but application is still running on <host>.**

The Identify Network Traffic Redirection scripts could be returning incorrect values. Make sure that the script on the Active host is returning 0 and the standby host is returning non-zero.

**When starting the scenario, I receive the error, ERROR: Unable to run scenario.**

The Identify Network Traffic Redirection scripts could be returning identical values. Make sure that the script on the Active host is returning 0 and the standby host is returning non-zero.

**When starting the scenario, I receive the warning, WARNING: Unable to connect to <Master or Replica> host<host> and error, ERROR: No network information for <Master or Replica> to check against.**

The Identify Network Traffic Redirection script on the Master or Replica could not be found by the engine.

# Chapter 19: Configuring the RHA NAT Utility for Various Network Setups

Use the RHA NAT utility when you have a network setup that uses firewalls, proxy servers, or a combination of both. These kinds of network setups restrict access to certain hosts and do not allow you to create scenarios. The following examples describe how to configure the NAT utility in various network setups.

This section contains the following topics:

# Example 1: Masters are behind a Closed Firewall

In this example, the Control Service (CS) and appliance cannot access the master. The following description explains the network setup.

- Masters are in a private LAN network that is behind a firewall (Firewall 1) without any port forwarding.

- Hosts in LAN 1 can access the public WAN through Firewall 1.

- The Control Service and replica are in a LAN network that is behind a firewall (Firewall2) with the port forwarding enabled.

- Hosts in LAN 2 can access the public WAN through Firewall 2. The following port forwarding is configured for Firewall 2.

  – Port 24000 is forwarded to the Control Service host.

  – Port 25000 is forwarded to the appliance.

- Both these private LAN networks are connected through a public WAN.

The following diagram illustrates the network setup.
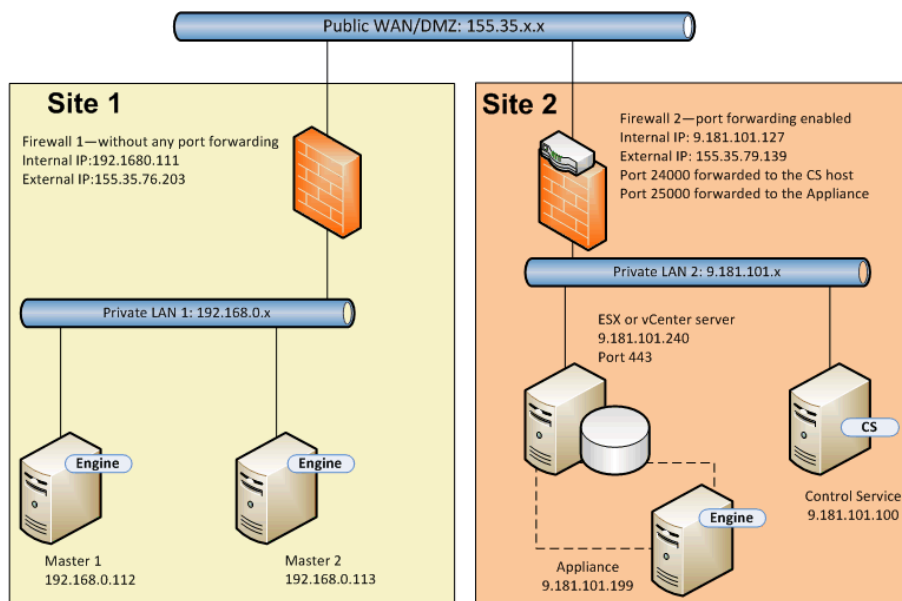


Let us see what happens if you create a Full System HA scenario without the NAT utility.

1. The Control Service cannot access the master in LAN1 because Firewall 1 has no port forwarding. You cannot create a scenario when the Control Service cannot access the master.

2.  When you run a backward scenario, data is copied from the appliance or switchover VM at Site 2 to the master at Site 1. This data replication is not possible because the Control Service cannot access Master 1.

To create a scenario for this setup, use the RHA NAT utility on Master 1.

The following tasks describe how to configure the NAT utility on the master while creating scenarios and performing BMR.

- ▪ Create a Full System HA Scenario
- ▪ Perform BMR from a Rewind Point
- ▪ Perform BMR using Reverse Replication

# Create a Full System HA Scenario

Before you create the Full System HA scenario, perform the following steps on Master 1.

1.  Run the RHA NAT utility from the engine folder.

2.  Enter the external IP, 155.35.79.139, and port numbers of Firewall 2.

    **Note:** The Control Service and appliance at Site 2 are forwarded from the same external IP (155.35.79.139) with different ports. So you have to add entries for each port as shown:

    - 155.35.79.139, 24000
    - 155.35.79.139, 25000
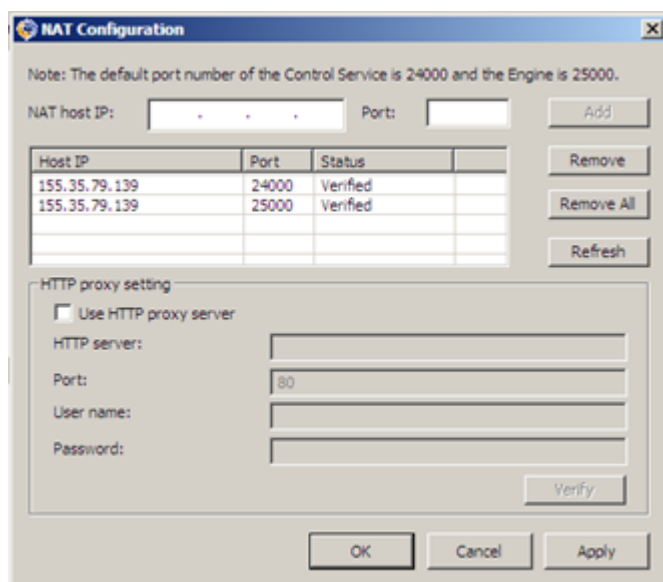


3.  Click OK to apply the configuration.

4.  Next, open the RHA Manager and click New on the standard toolbar to create the scenario.

5.  Select Full System with the product type HA and AR.

6.  Enter the IP address (192.168.0.112) and the default port number (25000) of the master.

7.  Enter the IP address (9.181.101.240) of the virtual platform and the default port number (443).

8.  Enter the IP address (155.35.79.139) of the appliance and the port number (25000). This port has a port forwarding to appliance (9.181.101.199).

**Note:** If you select Verify Arcserve RHA> Engine on Hosts, you get a message that RHA engine is not installed on 155.35.79.139. You get this error message because the verification is processed on Firewall 2.

9. Enter details of volume, resource pool, storage, scenario properties, and host properties similar to that of a full system scenario.

10. Select the switchover type and run the scenario.

# Perform BMR from a Rewind Point

The bare metal recovery lets you recover data and applications from a bookmark rewind point to a bare metal machine. Start the bare metal machine and configure the NAT utility.

**Follow these steps:**

1. Stop the scenario from the RHA manager.

2. Start Master 2 (bare metal machine) using the RHA BMR DVD or ISO image.

3. Verify the network adapter settings. If the settings are not correct, configure them manually.

4. Open the command prompt and navigate to the rha folder. Run the following command to open the RHA NAT utility.

   natutil

5. Run the following command to add Firewall 2 external IP address and port details.

   nat.addhost 155.35.79.139 24000

   nat.addhost 155.35.79.139 25000

   Port 24000 for Control Service and 25000 for the appliance.

6. Run the following command to apply the settings and establish a connection between Master 2 and the Control Service.

   nat.apply

7. Next, open the RHA Manager and click New on the standard toolbar to create the scenario.

8. Click Restore Data to open the Data Restore Wizard.

9. Select Recovery Type (Bare Metal Recovery) and the recovery point.

10. Enter the following IP address and port details on the Recovery Destination page.

    **Destination IP:** 192.168.0.113, 25000

11. Click Verify to verify the destination host.

12. Enter details of volume, resource pool, storage, scenario properties, and host properties similar to that of a full system scenario.

13. Click Run to start the recovery scenario.

# Perform BMR Using Reverse Replication

You can recover application and data after failover using reverse replication for this network setup. After switchover or failover occurs, a Run button is available to launch the reverse replication process. A new switchover VM is created. This VM contains the same data as the master but has a different IP address (9.181.101.152 instead of 192.168.0.112).

In this example, the switchover is between the switchover VM at Site 2 and the bare metal machine (Master 2) at Site1. To make this switchover VM accessible, you set a new port forwarding (25001) on Firewall 2.
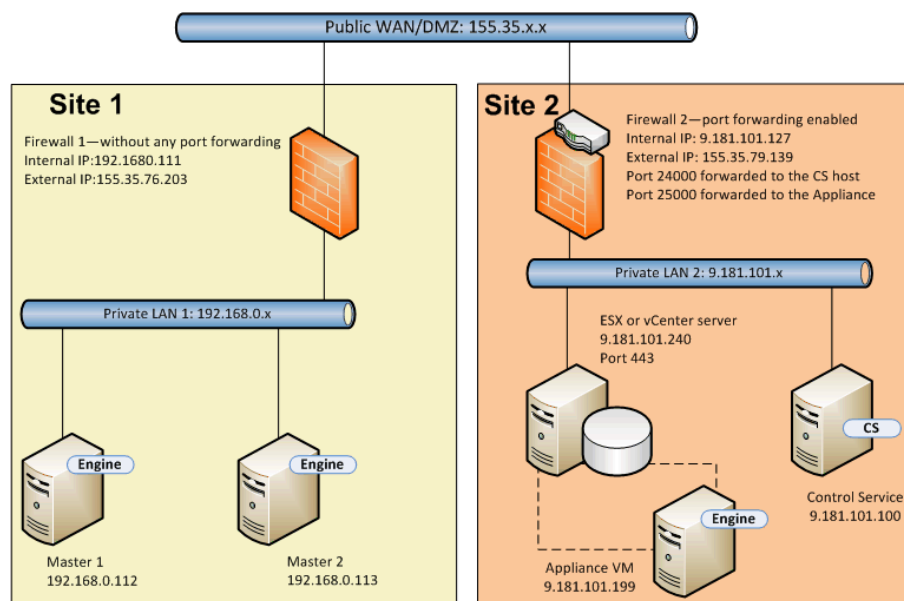


Now you have the Control Service, appliance, and a new switchover VM. Enter the details of all three machines in the NAT utility on the master.

**Follow these steps:**

1.  Start the bare metal machine (Master 2) using the BMR DVD or ISO image.

2.  Verify the network adapter settings. If the settings are not correct, configure them manually.

3.  Open the command prompt and navigate to the rha folder. Run the following command to open the RHA NAT utility.

    natutil

4.  Run the following command to add Firewall 2 external IP address and port details.

    nat.addhost 155.35.79.139 24000

nat.addhost 155.35.79.139 25000

nat.addhost 155.35.79.139 25001

Port 24000 for the Control Service, 25000 for the appliance, and 25001 for the switchover VM.

5.  Run the following command to apply the settings and establish a connection between the master, Control Service and appliance.

nat.apply

6.  Next, open the RHA Manager to create the backward scenario. Select the scenario that performed switchover and click Run to open the Data Restore Wizard.

7.  Enter the following IP address and port details on the Reverse Replication page.

**Source Name/IP:** 155.35.79.139, 25001

**Destination IP:** 192.168.0.113, 25000

8.  Click Verify to verify the destination host.

9.  Enter details of volume, resource pool, storage, scenario properties, and host properties similar to that of a full system scenario.

10. Click Run to start the backward scenario and run the reverse replication pro-cess.

After the synchronization is complete, start the manual switchback to make the bare metal machine (Master 2) live.

# Example 2: Masters are behind a Closed Firewall that uses a Proxy Server

This example is similar to Example 1, the only difference is that there is an additional proxy server at Site 1. The following description explains the network setup.

- Masters are in a private LAN network that is behind a firewall (Firewall 1) without any port forwarding.

- Firewall 1 is connected to a proxy server (155.35.76.123).

- Hosts in LAN 1 can access the public WAN through Firewall 1.

- The Control Service and the replica are in a LAN network that is behind a firewall (Firewall2) with the port forwarding enabled.

- Hosts in LAN 2 can access the public WAN through Firewall 2. The following port forwarding is configured for Firewall 2.

    - Port 24000 is forwarded to the Control Service host.

    - Port 25000 is forwarded to the appliance.

- Both these private LAN networks are connected through a public WAN.

The following diagram illustrates the network setup.
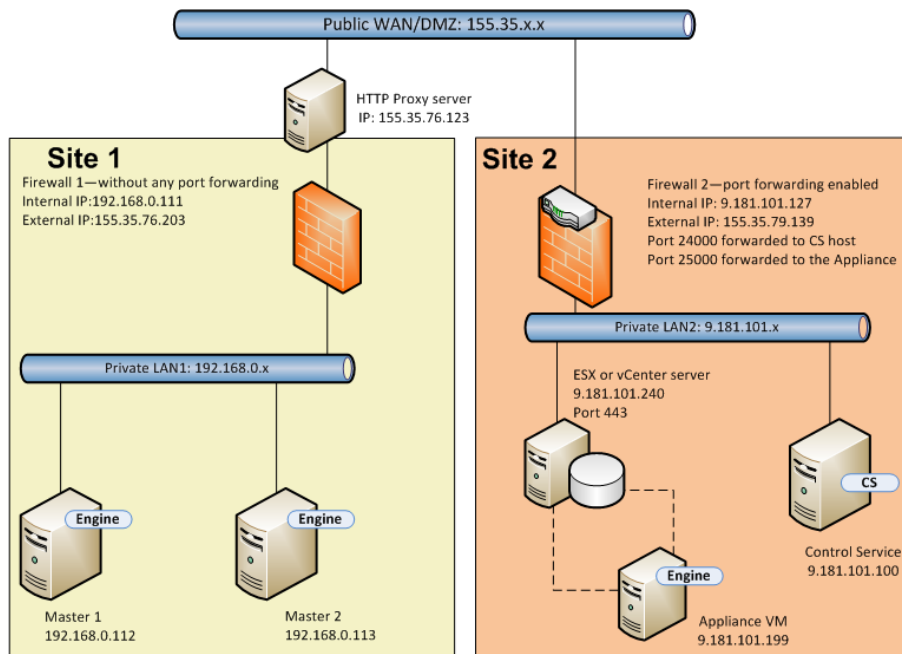


This example is similar to Example 1. You only need to perform the following additional steps.

1. On the master, specify the following proxy server details in the NAT utility.

    - **HTTP Server:** 155.35.76.123

    - **Port:** 80

    - **User name:** <domain\username>

    - **Password:** <password>
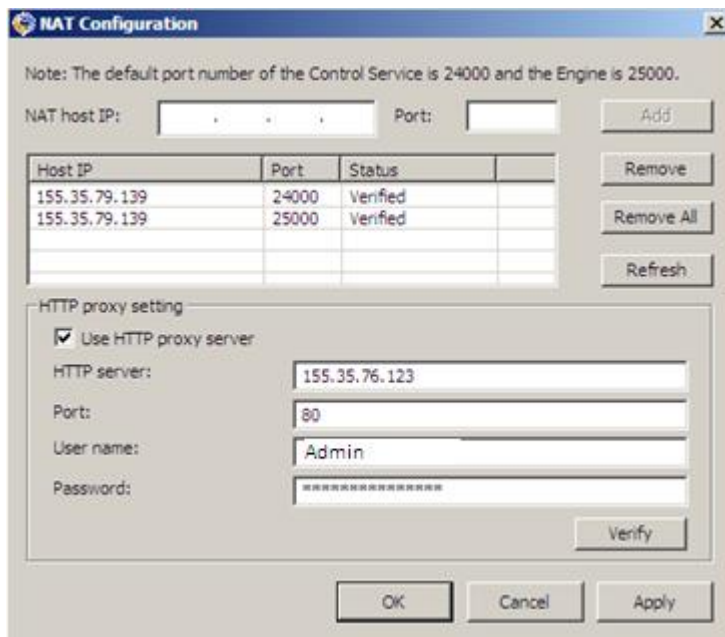


2. While you create the Full System scenario, specify the following host properties in the Master and Replica Properties screen.

    - Set the HTTP Proxy Server Configuration to ON.

    - Enter the proxy server details.
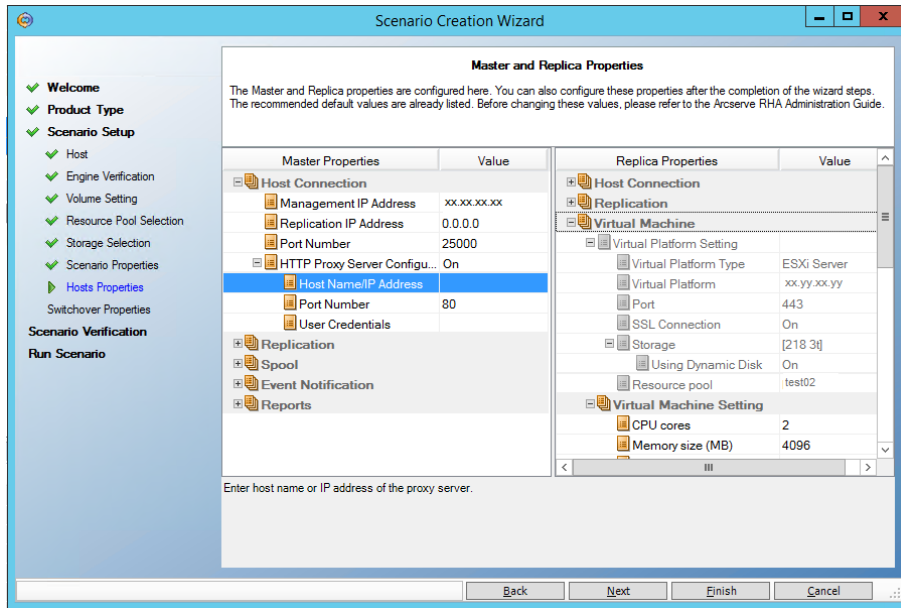
# Example 3: Appliance and CS are Behind a Closed Firewall

In this example, the master cannot access the replica. The Control Service and replica are behind a firewall without any port forwarding. The following description explains the network setup.

- Masters are in a private LAN network (LAN 1) that is behind a firewall (Firewall 1) with the port forwarding enabled.

- Hosts in LAN 1 can access the public WAN through Firewall1. The following port forwarding is configured for Firewall 1.

  - Port 25000 is forwarded to Master 1.

  - Port 25001 is forwarded to Master 2.

- The Control Service and replica are in a private LAN network (LAN 2) that is behind a firewall (Firewall2) without any port forwarding.

- Both these private LAN networks are connected through a public WAN.

The following diagram illustrated the network setup.
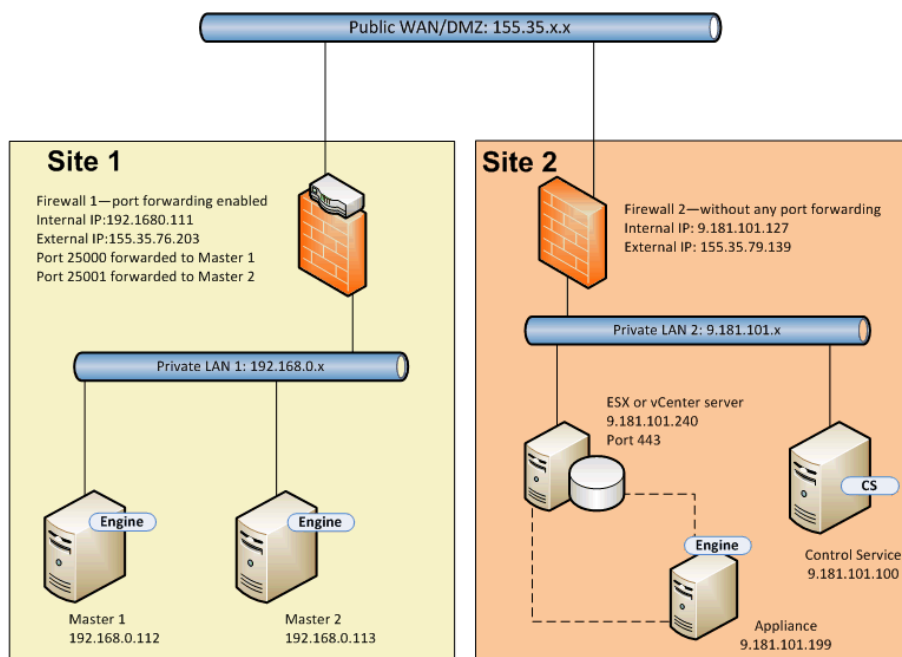


In this setup, the Control Service and the appliance can access the master but the master cannot access the appliance. For a forward scenario, the master cannot copy data to the replica.

To create a scenario for this kind of setup, configure the RHA NAT utility on the appliance.

The following tasks describe how to configure the NAT utility on the appliance while creating a scenario and performing BMR.

- Create a Full System HA Scenario

- Perform BMR from a Rewind Point

- Perform BMR using Reverse Replication

# Create a Full System HA Scenario

Before you create the Full System HA scenario, perform the following steps on the appliance.

**Follow these steps:**

1.  Launch the RHA NAT utility from the engine folder on the appliance.

2.  Enter the following IP addresses and port details.

    155.35.76.203, 25000

3.  Click OK to apply the configuration.

4.  Next, open the RHA Manager and click New on the standard toolbar to create the Full System scenario.

5.  Select Full System with the product type HA and AR.

6.  Enter the Master IP address (155.35.76.203) and port number (25000). This IP is the external IP address of Firewall 1 with port 25000 forwarded to Master 1 (192.168.0.112).

7.  Enter the IP address (9.181.101.240) of the virtual platform and the default port number (443).

8.  Enter the IP address (9.181.101.199) of the appliance and the port number (25000).

9.  Enter details of volume, resource pool, storage, scenario properties, and host properties similar to that of a full system scenario.

10. Select the switchover type and run the scenario.

## Perform BMR from a Rewind Point

The steps are almost the same as in Example 1, Perform BMR from a Rewind Point. Only the following two steps are different from Example 1.

- After the bare metal machine starts up, you do not need to configure the NAT utility on the master. The master can access the appliance using Firewall 1 and port forwarding.

- On the Recovery Destination page, enter the following external IP address of Firewall 1 and the forwarded port number.

  **Destination IP:** 155.35.76.203, 25000

# Perform BMR Using Reverse Replication

When switchover occurs, the switchover VM appears at Site 2. This switchover VM can connect to the standby host (Master 2) by using the Firewall 1 external IP address and port forwarding. So you do not need to use the NAT utility on the switchover VM.

# Example 4: CS is on a Public WAN

This example is similar to Example 3 except that the CS host is on the public WAN with a public IP address (155.35.76.123). The following description explains the network setup.

- Masters are in a private LAN network (LAN 1) that is behind a firewall (Firewall 1) with the port forwarding enabled.

- Hosts in LAN 1 can access the public WAN through Firewall 1. The following port forwarding is configured for Firewall 1.

  - Port 25000 is forwarded to Master 1.

  - Port 25001 is forwarded to Master 2.

- The Control Service is on the public WAN.

- The appliance is in a private LAN network (LAN 2) that is behind a firewall (Firewall2) without any port forwarding.

- Both these private LAN networks are connected through a public WAN.

The following diagram illustrates the network setup.
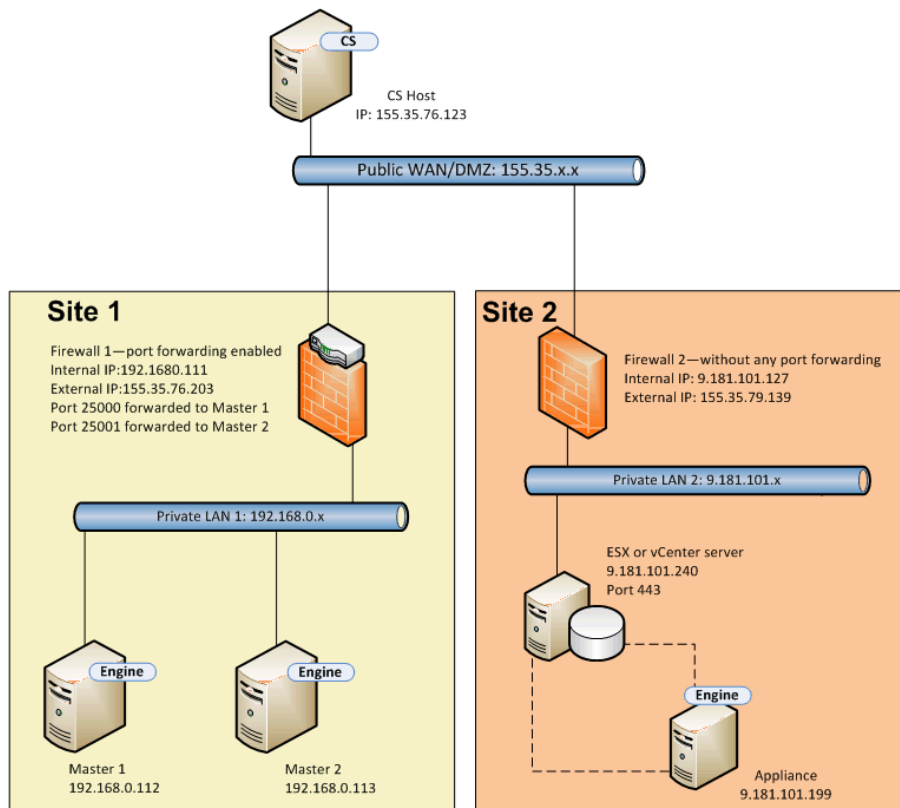
The Control Service can connect to the master but cannot connect to the appliance. So configure the NAT utility on the appliance to establish a connection between the Control Service and the master. The other steps are same as in Example 3.

# Example 5: Masters, Appliance, and CS are Behind Port Forwarded Firewalls

In this example, the network setup is such that in both sites the hosts are behind firewalls that have port forwarding enabled.

The following description explains the network setup.

- Masters are in a LAN network (LAN 1) that is behind a firewall (Firewall 1).

- Port forwarding is enabled for Firewall 1.

- The Control Service (CS) and replica are in a LAN network (LAN 2) that is behind a firewall (Firewall2).

- Port forwarding is enabled for Firewall 2.

- Both these LAN networks are connected using a public WAN.

The following diagram describes the network setup.
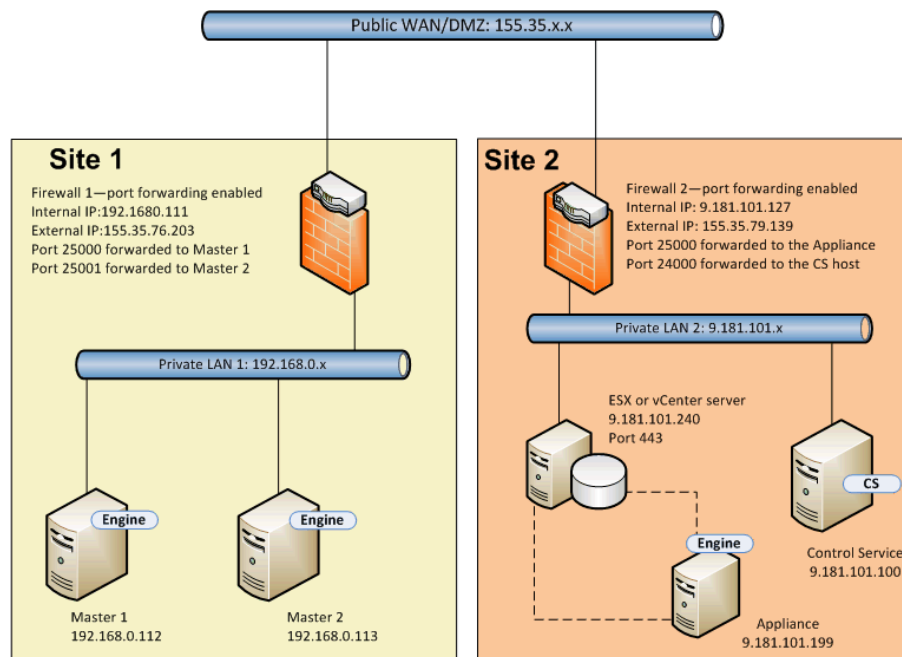


Port forwarding is enabled for both firewalls, so masters can access the appliance and Control Service. Similarly the Control Service and the replica can access the masters. For this setup, you can create a scenario in the following two ways.

- **Using the NAT utility:** Use the NAT utility either on masters or on the appliance to create an HA scenario, see either Example 1 or Example 2.

- **Without using the NAT utility:** Use the external IP addresses of Firewall 1 and Firewall 2 with port forwarding details to create an HA scenario.

# Chapter 20: Enable Mutual Authentication

To prevent replays and session injection or login issues between RHA roles (Control Service and Engines), in r16.5 Service Pack 7 RHA introduced a tool to enable mutual authentication and maintain the certificates and private keys (with or without encryption).

The mutual authentication is disabled by default. If enabled, any remote connection is verified on the local server. If the certificate of the client does not exist in trust store on the local server, the connection is refused.

The certutl.exe tool is located under Control Service and Engine installation folder and configures the trusted certificates store on Control Service and Engine. Using this tool, the user has following benefits:

- Set to enable or disable SSL certificate verification before establishing connections.

- Set the SSL certificate and unencrypted or encrypted private key used by local server.

- Add new certificate to local trusted certificate store.

- Revoke certificate from local trusted certificate store.

- Add or Remove the revocation list for certificates.

- Test the possibility to establish THE SSL connection to remote server.

- Encrypt all into encrypted certificate store with unique encryption key to prevent from copying local certificate store to another host.

**Notes:**

- The updates to the configuration of mutual authentication will only function by restarting the services of Control Service or Engine.

- To initiate the certificate store, use the built-in administrator account to run certutl.exe tool.

- To use SSL Certification Verification, Control service and all engines need r16.5 RHA Service Pack 7. If Control Service has Service Pack 7 while the engines have lower versions, run scenarios will fail because low versions of RHA do not support SSL Certification.

This section contains the following topics:

# List All Available Commands and Display Current Configuration

As the first step of enabling mutual authentication, you need to list all available commands and display current configuration.

**Follow these steps:**

1. Using Command Prompt, navigate to the installation directory of Control Service (CS) or Engine.

2. Run certutl.exe.

3. In the command shell, use certutl.l to list all commands.
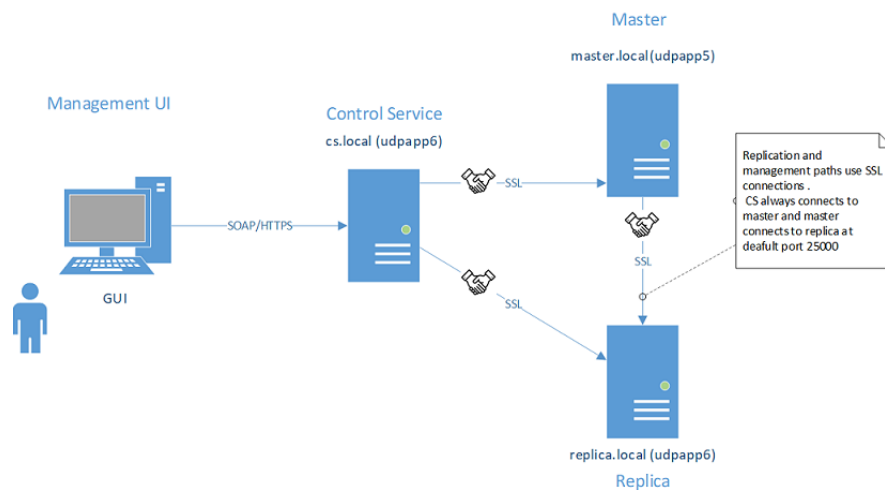


4. Use *q* to quit the command shell.

5. Use certutl.display to display the current configuration in the certificate store, including the certificate and private key currently used on the host, SSL certificate verification (is enabled or not), and all certificates added in the local trust certificate store.



6. Use command with parameter *-h* to display the command help.

7.  Use certutl.testsslconn to test if you can establish SSL connection on the remote server.

# Enable or Disable SSL Certification Verification

Use certutl.enablesslconnverify to enable or disable SSL certificate verification before establishing the connections.

Default: *disabled*

Usage: certutl.enablesslconnverify True/False

# Set / Reset SSL Certificate and Private Key

You need to set or reset SSL certificate and private key.

**Follow these steps:**

1. Use certutl.setcertpkey to set the SSL certificate and private key used by local server.

   Supporting unencrypted or encrypted private key, the key validates the certificate, checks the private key or the password for the encrypted private key, and then checks if the private key matches the certificate.1. Use certutl.setcertpkey to set the SSL certificate and private key used by local server. Supporting unencrypted or encrypted private key, the key validates the certificate, checks the private key or the password for the encrypted private key, and then checks if the private key matches the certificate.

   **Usage:** certutl.setcertpkey <SSL certificate file> <RSA private key file> [password for encrypted private key]



2. Use certutl.resetcertpkey to clear the SSL certificate and private key that is set by certutl.setcertpkey.

# Add / Revoke Certificates to the Local Trust Certificate Store

You can add or revoke certificates to the local trust certificate store.

**Follow these steps:**

1. Use certutl.addtrustcert to add new certificates to the local trust certificate store.

2. Use certutl.revoketrustcert or certutl.revoketrustcertbyfingerprint to revoke the certificate from the local trust certificate store.

**Usage:**

certutl.addtrustcert <SSL certificate file>



certutl.revoketrustcert <SSL certificate file>

certutl.revoketrustcertbyfingerprint <certificate SHA1 fingerprint>

# Set / Reset the URL for CRL Revocation List

You need to set or reset URL for CRL revocation list.

**Follow these steps:**

1. Use certutl.setURLforCRLrevocationlist command to set the URL of CRL revocation list.

   **Usage:** Certutl. setURLforCRLrevocationlist <URL of CRL revocation list>

2. Use certutl.resetURLforCRLrevocationlist to reset the URL of CRL revocation list.

# Example to Configure Mutual Authentication

The example scenario in this section explains how to enable mutual (certificates-based) authentication for hosts participating in Arcserve Replication and High Availability scenario. For example, let us consider simple File Server scenario where data gets replicated from host A to host B. The example has mutual authentication enabled between all Arcserve Replication and High Availability roles participating in the scenario that includes Control Service, Master engine, and Replica engine.
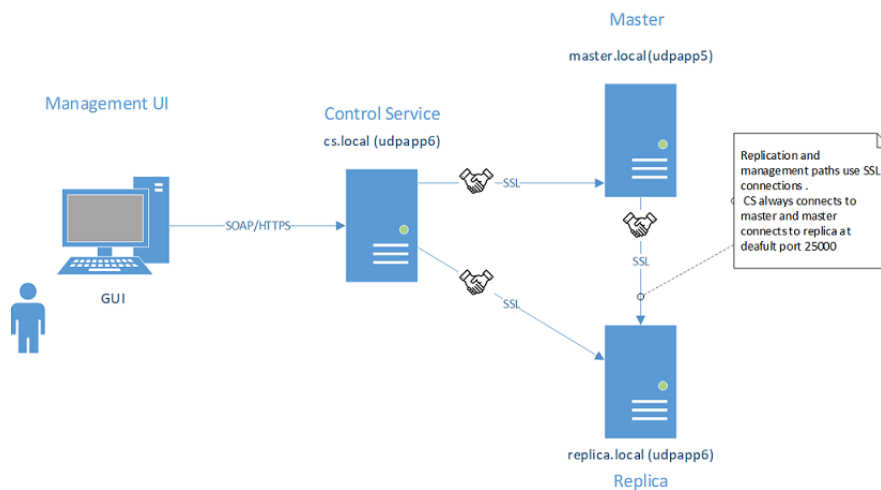


For two hosts to mutually trust each other: Each host must use a valid SSL certificate for establishing management connections, remote peer's certificate should be added into host's trusted certificates stores and mutual authentication option should be enabled on each host.

This example has three roles / hosts: Control service (CS), master, and replica. To trust each other, all the three roles need the following:

1. Generate valid SSL certificates for CS, master, and replica.

2. Configure corresponding role to use the SSL certificate for management connections.

3. On each host (role) enable mutual authentication feature and add public certificates of all trusted hosts into role's trusted certificate store.

**Notes:**

- For using commercial certificates or own certification authority (CA) to generate certificates, add all certificates from chain of authority into trusted certificate store of role. To verify certificate signed by other certificate / CA, the

trusted store of role needs to have public certification of CA / signing certificate.

- For importing several certificates by the PEM file, make one PEM file that only contains just one certificate and import those PEM files one by one.

- All examples in this guide are also valid for UNIX/Linux hosts. For Unix/Linux, use /opt/ARCserve/bin/certutl utility to configure SSL settings for RHA engine. Configuration steps are also similar to those in this example. To restart engine on Linux, use the following command: *service ws_rep restart*.

The example is explained with the following topics:

- Creating Self-Signed Certificates
- Configure Roles to use SSL Certification
- Test SSL Connectivity among three roles

## Creating Self-Signed Certification

The example uses self-signed certificates. You may also use commercial certificates and/or use own certification authority to create chain of trust.

To generate self-signed certificates, use OpenSSL utility as displayed in the screenshot below. This example is running OpenSSL from bash for Windows. To install OpenSSL, use command *sudo apt-get install openssl*.)

Now, the requirement is to generate private key and public certificate for CS , master, and replica. To generated, run command thrice changing names of output files.

Generating CS key / certificate pair:



Generating Master key / certificate pair:



Generating Replica key / certificate pair:

Three pairs of private keys and public certificates are generated.



The last step involves testing the SSL connectivity among rules.

## Configuring SSL Certificate for Three Roles

The next step in example is to configure use of generated SSL certificates for every role, add peers public certificates into local trusted store and enable mutual authentication on corresponding hosts. View the following links to configure roles:

- Configure Control Service
- Configure Master Engine
- Configure Replica Engine
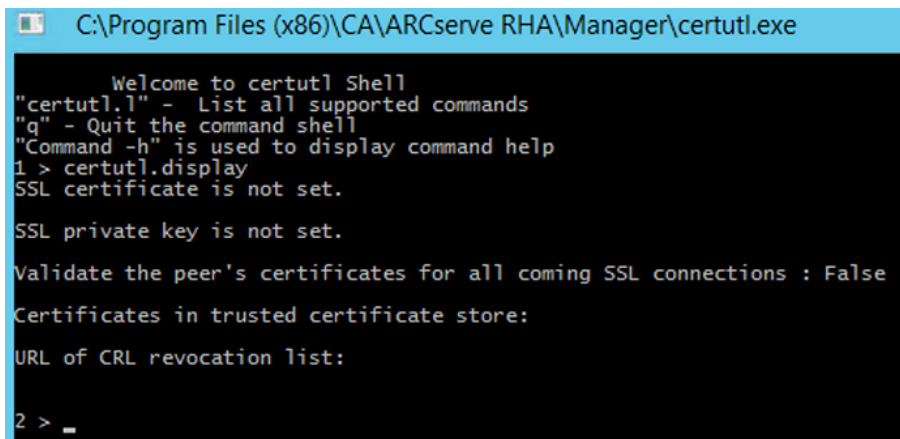
### Configure Control Service

On Control Service host, change directory to CS installation directory and start certutil.exe.

Type the following command to display the current SSL configuration:

*certutil.display*

The screenshot below displays that SSL was not previously configured for this host:
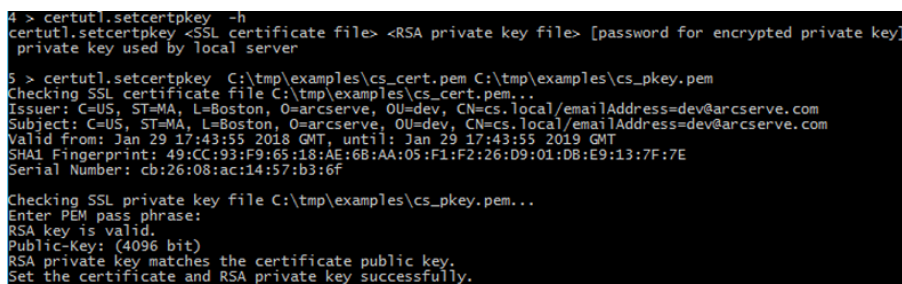


To generate self-signed certificates, use OpenSSL utility as displayed in the screenshot below. This example is running OpenSSL from bash for Windows. To install OpenSSL, issue command "sudo apt-get install openssl".)

**Follow these steps::**

1. Set SSL certificate of CS for management connections.



   The above command sets the previously generated private key and public certificate to be used by CS for all management connections (generated certificate and key were copied to CS host before issuing command).

2. Add public certificates of master and replica (peers) to local trusted certificate store on CS.

```
6 > certutl.addtrustcert -h
certutl.addtrustcert <SSL certificate file> Add new certificate to local certificate store

7 > certutl.addtrustcert C:\tmp\examples\master_cert.pem          add master's cert to local store
Checking SSL certificate file C:\tmp\examples\master_cert.pem...
Issuer: C=US, ST=MA, L=bos, O=dev, OU=arcserve, CN=master.local/emailAddress=dev@arcserve.com
Subject: C=US, ST=MA, L=bos, O=dev, OU=arcserve, CN=master.local/emailAddress=dev@arcserve.com
Valid from: Jan 29 17:53:15 2018 GMT, until: Jan 29 17:53:15 2019 GMT
SHA1 Fingerprint: 54:79:45:45:34:90:CB:12:E9:AC:24:FE:A6:40:10:AC:EE:8D:74:BA
Serial Number: 90:d8:ed:c8:61:10:d1:56

Add SSL certificate to local certificate store successfully.

8 > certutl.addtrustcert C:\tmp\examples\replica_cert.pem         add replica's cert to local store
Checking SSL certificate file C:\tmp\examples\replica_cert.pem...
Issuer: C=US, ST=ma, L=bos, O=arcserve, OU=dev, CN=replica.local/emailAddress=dev@arcserve.com
Subject: C=US, ST=ma, L=bos, O=arcserve, OU=dev, CN=replica.local/emailAddress=dev@arcserve.com
Valid from: Jan 29 17:54:34 2018 GMT, until: Jan 29 17:54:34 2019 GMT
SHA1 Fingerprint: 17:E0:EB:EE:C5:BE:66:D6:36:53:10:CB:E2:AE:50:CB:64:C6:37:6F
Serial Number: fa:8a:cc:dd:08:4d:d6:16

Add SSL certificate to local certificate store successfully.
```

3. Enable verification of SSL certificates for all incoming connections.

   **Note:** This option must be enabled to ensure that CS accepts connections only from hosts whose public certificates were added into local trusted cert store.

```
12 > certutl.enableSSLconnverify -h
certutl.enableSSLconnverify <True or False> Validate the peer's certificates for all SSL connections or not

13 > certutl.enableSSLconnverify True
Enable to validate the peer's certificates for all SSL connections successfully.
```

4. Verify if SSL configuration for CS look similar to the screenshot below:

```
14 > certutl.display
SSL certificate:
Issuer: C=US, ST=MA, L=Boston, O=arcserve, OU=dev, CN=cs.local/emailAddress=dev@arcserve.com      CS uses SSL cert for management
Subject: C=US, ST=MA, L=Boston, O=arcserve, OU=dev, CN=cs.local/emailAddress=dev@arcserve.com     connections
Valid from: Jan 29 17:43:55 2018 GMT, until: Jan 29 17:43:55 2019 GMT
SHA1 Fingerprint: 49:CC:93:F9:65:18:AE:6B:AA:05:F1:F2:26:D9:01:DB:E9:13:7F:7E
Serial Number: cb:26:08:ac:14:57:b3:6f

SSL private key:
RSA key is valid.
Public-Key: (4096 bit)                                          CS will accept connections only fro hosts whose certs
                                                                are in trusted cert store
Validate the peer's certificates for all coming SSL connections : True

Certificates in trusted certificate store:
Certificate #0
Issuer: C=US, ST=MA, L=bos, O=dev, OU=arcserve, CN=master.local/emailAddress=dev@arcserve.com
Subject: C=US, ST=MA, L=bos, O=dev, OU=arcserve, CN=master.local/emailAddress=dev@arcserve.com    CS trusts to master's and replica's
Valid from: Jan 29 17:53:15 2018 GMT, until: Jan 29 17:53:15 2019 GMT                              certs
SHA1 Fingerprint: 54:79:45:45:34:90:CB:12:E9:AC:24:FE:A6:40:10:AC:EE:8D:74:BA
Serial Number: 90:d8:ed:c8:61:10:d1:56
Certificate #1
Issuer: C=US, ST=ma, L=bos, O=arcserve, OU=dev, CN=replica.local/emailAddress=dev@arcserve.com
Subject: C=US, ST=ma, L=bos, O=arcserve, OU=dev, CN=replica.local/emailAddress=dev@arcserve.com
Valid from: Jan 29 17:54:34 2018 GMT, until: Jan 29 17:54:34 2019 GMT
SHA1 Fingerprint: 17:E0:EB:EE:C5:BE:66:D6:36:53:10:CB:E2:AE:50:CB:64:C6:37:6F
Serial Number: fa:8a:cc:dd:08:4d:d6:16

URL of CRL revocation list:
```

5. To complete configuration on CS, perform the following steps:

   a. Close certutil.exe by typing *q*.

   b. Restart Control Service by using commands: *sc stop CAARCserveRHAManager* and *sc start CAARCserveRHAManager*.

### Configure Master Engine

On master host, all configuration steps are like those performed on CS with the following exceptions:

- Start certutil from the installation directory of engine on master.

  **Default:** C:\Program Files\CA\ARCserve RHA\Engine\certutl.exe

- In certutil.exe, use master_cert.pem and master_pkey.pem files in the step that involves setting the SSL certificate for management/replication connections.

- Add public certificates of CS and replica to master's trusted certificate store.

The following screenshot displays the capturing configuration process on master engine:
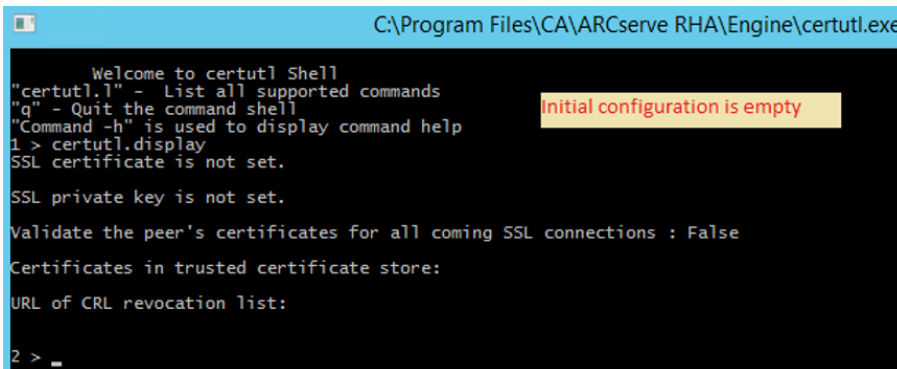


**Follow these steps::**

1. Performing configuration as displayed in the screenshot below.
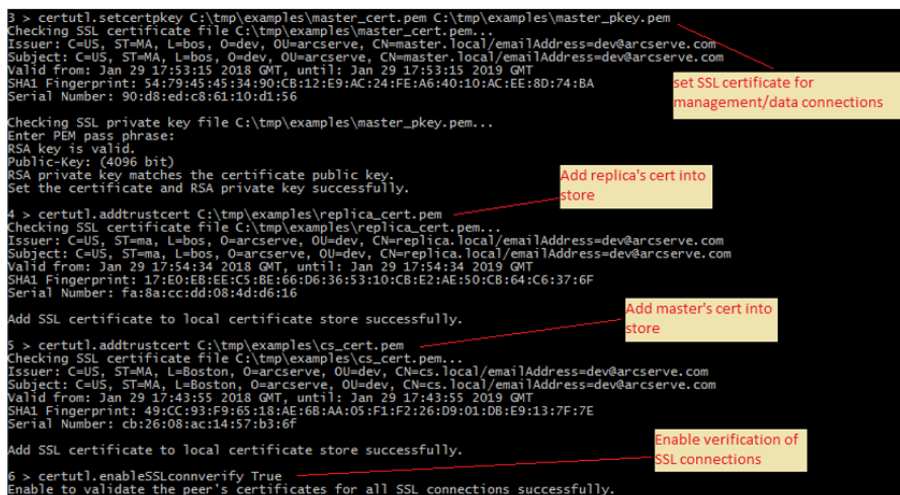


2. Verify if the Master's configuration appears as displayed in the screenshot below.

```
7 > certutl.display
SSL certificate:
Issuer: C=US, ST=MA, L=bos, O=dev, OU=arcserve, CN=master.local/emailAddress=dev@arcserve.com
Subject: C=US, ST=MA, L=bos, O=dev, OU=arcserve, CN=master.local/emailAddress=dev@arcserve.com
Valid from: Jan 29 17:53:15 2018 GMT, until: Jan 29 17:53:15 2019 GMT
SHA1 Fingerprint: 54:79:45:45:34:90:CB:12:E9:AC:24:FE:A6:40:10:AC:EE:8D:74:BA
Serial Number: 90:d8:ed:c8:61:10:d1:56

SSL private key:
RSA key is valid.
Public-Key: (4096 bit)

Validate the peer's certificates for all coming SSL connections : True

Certificates in trusted certificate store:
Certificate #0
Issuer: C=US, ST=ma, L=bos, O=arcserve, OU=dev, CN=replica.local/emailAddress=dev@arcserve.com
Subject: C=US, ST=ma, L=bos, O=arcserve, OU=dev, CN=replica.local/emailAddress=dev@arcserve.com
Valid from: Jan 29 17:54:34 2018 GMT, until: Jan 29 17:54:34 2019 GMT
SHA1 Fingerprint: 17:E0:EB:EE:C5:BE:66:D6:36:53:10:CB:E2:AE:50:CB:64:C6:37:6F
Serial Number: fa:8a:cc:dd:08:4d:d6:16
Certificate #1
Issuer: C=US, ST=MA, L=Boston, O=arcserve, OU=dev, CN=cs.local/emailAddress=dev@arcserve.com
Subject: C=US, ST=MA, L=Boston, O=arcserve, OU=dev, CN=cs.local/emailAddress=dev@arcserve.com
Valid from: Jan 29 17:43:55 2018 GMT, until: Jan 29 17:43:55 2019 GMT
SHA1 Fingerprint: 49:CC:93:F9:65:18:AE:6B:AA:05:F1:F2:26:D9:01:DB:E9:13:7F:7E
Serial Number: cb:26:08:ac:14:57:b3:6f

URL of CRL revocation list:
```
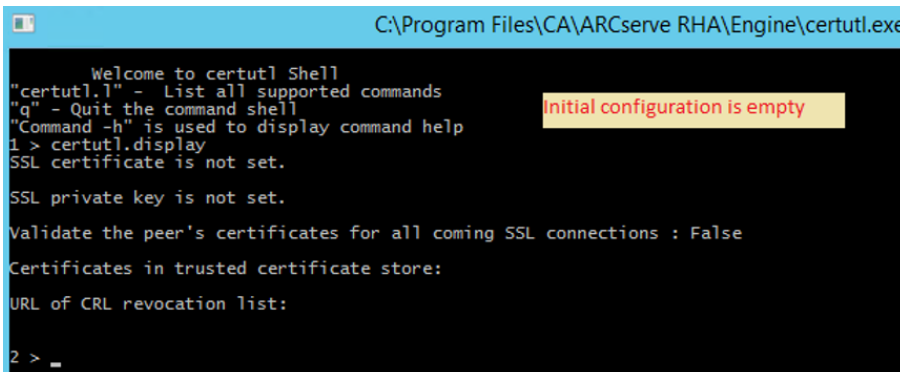
3. To complete configuration on Replica, perform the following steps:

   a. Close certutil.exe by typing *q*.

   b. Restart Master Engine Service by using commands: *sc stop CAARCserveRHAEngine* and *sc start CAARCserveRHAEngine*.

**Configure Replica Engine**

On Replica host, all configuration steps are like those performed on CS with the following exceptions:

- Start certutil from the installation directory of engine on replica host.

  **Default:** C:\Program Files\CA\ARCserve RHA\Engine\certutl.exe

- In certutil.exe, use replica_cert.pem and replica_pkey.pem files in the step that involves setting the SSL certificate for management/replication connections.

- Add public certificates of CS and replica to trusted certificate store of replica.

The following screenshot displays the capturing configuration process on replica engine:



**Follow these steps::**

1. Performing configuration as displayed in the screenshot below.

2. Verify if the replica configuration appears as displayed in the screenshot below.



3. To complete configuration on Master, perform the following steps:

   a. Close certutil.exe by typing *q*.

   b. Restart Master Engine Service by using commands: *sc stop CAARCserveRHAEngine* and *sc start CAARCserveRHAEngine*.

## Testing SSL Connectivity among Roles

Use certutil.exe to verify whether mutual trust among all roles participating in the Arcserve Replication and High Availability scenario was configured properly.

Execute the following on CS:



You may perform similar test checking connectivity from replica to master and CS and from master to CS and replica (using certutil.testSSLconn command).

# Chapter 21: Arcserve RHA Troubleshooting

The following section provides you information on some of the error messages that you may receive - when they occur, what is their meaning, and how to resolve the problems that caused them.

This section contains the following topics:

## Troubleshooting Tips

- If you trigger a manual switchover, Arcserve RHA sends the last journal event to all replica nodes before stopping the scenario. When the Master is up, initial synchronization is skipped. If you wish to force resynchronization from the owner to other nodes, click the option, Force a data synchronization.

- Whenever you change a node in the cluster (for example, eject, add), you should rescan the cluster nodes for this resource. To rescan cluster nodes, click the option, Rescan cluster nodes for this resource.

## Spool Limit Exceeded

**CR00404 "Spool limit exceeded. The scenario is stopped"**

**Reason:**

This message can refer either to the spool on the Master or to the spool on the Replica. It means that the value you entered in one of the Spool properties - **Max Spool**

**Size** or **Min Disk Free Size** -  exceeded. As a result of reaching the threshold, the system issues an error and stops the ongoing replication.

Several factors can cause the spool growth:

1. On the Master and Replica - when a real-time Antivirus scanning is active, it filters every replicated file before applying the change to the data file. This filtering process causes the replicated files to cache in the spool on both the Master and Replica, before they are transferred or applied. As a result, a bottleneck appears, and the spool limit may exceed.

2. On the Master - when there is a low bandwidth connection between the Master and the Replica, and many updates occur on the Master, the spool limit may exceed.

3. On the Replica - when synchronization is initiated during an ongoing replication, the new updates from the Master are stored on the Replica spool. Only when the synchronization is finished, the replicated files that are cached in the spool directory are applied to the Replica data. If the synchronization is long, or many changes occurred on the Master in the meantime, the spool limit may exceed.

4. On the Replica - during the Assured Recovery test, the data changes that take place on the Master are sent to the Replica, but they are not immediately applied. Instead, these changes are accumulated and stored in the spool, and only when the testing is completed they are applied to the Replica data. This may cause the spool limit to exceed.

5. On the Replica - when there is a lot of activity on the data directory of the Replica server, a Disk IO queue may be formed. This causes the replicated files to begin cache in the spool, waiting for a sequential access to the disk subsystem, in order to apply the data changes to the files on the disk.

6. There is high Disk IO on the data directory on the Replica server. Disk IO is causing disk writes to queue, forcing the replication files to begin caching in the spool, waiting for sequential access to the disk subsystem in order to apply the data changes to the files on disk.

7. To resolve consider running Perform Counters to review Disk IO. consider moving spool to a volume which has relatively low Disk IO. It is not recommended to place the Spool on the same volume where the data files are located. In the case of an application server such as Exchange, SQL, Oracle, etc) the Spool should not be placed on a volume that contains either the DB files or the Transaction logs.

**Note**:

- When there is a disconnection between a Master and a Replica, the spool on the Master that stores changes for the disconnected Replica will be emptied.

- When the spool or the spool drive is full, Arcserve RHA stops the scenario. Once you restart the scenario, the spool is emptied, so the scenario can start running. However, if you do not make any change in the Spool definition or the Spool drive, this problem will occur again.

**Action:**

You can do one or all of the following actions:

- Exclude the Arcserve RHA Spool directory from AV scanning on all hosts involved in the replication.

- Decrease the value you entered in the **Spool** - **Min Disk Free Size** property.

- Increase the value you entered in the **Spool** - **Max Spool Size** property.

- Run Perform Counters to review the Disk IO activity. If necessary, change the **Spool Directory** location, and select a free and relatively inactive drive.

**Note**: It is not recommended to place the Arcserve RHA Spool on the same volume where the data files are located. In the case of database servers, such as Exchange, SQL, and Oracle, the Spool should not be placed on a volume that contains either the DB files or the Transaction logs.

- Delete or move files from the current spool drive, and reduce its other activities.

# The Disk is Full

**"CR01488" "The disk is full. Stopping replication"**

**Reason:**

This message can refer either to the Master or to the Replica. However, in most cases it refers to lack of free disk space on the Replica, which prevents the replication from continuing.

Several common factors can cause the Replica disk to become full:

1. When the size of the replicated data is bigger than the Replica disk size, the Replica disk will become full before all the replicated data is applied.

2. When the Replica contains other data, it may not leave enough free space to store the replicated data.

3. When VSS snapshots are created on a scheduled basis and are stored on the Replica, they may gradually occupy too much disk space.

**Action:**

You can do one or all of the following actions:

- Free space on the full disk for the replicated data.

- Check and compare the disk volume capacity of the Replica and the size of the replicated data, and select a Replica host with enough free disk space.

- If the system is configured to create VSS snapshots on a scheduled basis, you can either remove old snapshots or change the VSS snapshot schedule and properties.

## EM03100

**Error EM03100 Replica %1 doesn't join into the deployment of master %2**

**Applies to Dynamics CRM**

**Action:**

Re-install CRM on the Replica server and choose the option, Connect to existing deployment.

## EM03101

**Error EM03101 Roles (except SQL Server) installed on Replica are not same on Master %1 %2**

**Applies to Dynamics CRM**

**Action:**

Install the same roles on both the Master and Replica.

## EM03102

**Error EM03102 Role %1 is absent.**

**Applies to Dynamics CRM**

**Action:**

Install the same roles on both Master and Replica servers.

## EM03103

**Error EM03103 Role %1 is redundant**

**Applies to Dynamics CRM**

**Action:**

Install the same roles on both the Master and Replica servers.

# Renew an Expired SSL Certificate

When you are using an SSL Certificate to secure communication, you may need to renew expired certificates or install new ones. The following procedure applies to both authorized or self-signed SSL certificates. Contact Support to obtain a new certificate.

**To renew an expired SSL certificate**

1. Obtain a new certificate and install it on the machine where the Control Service is running.

   **Note**: You do not need to stop the Control Service during this procedure.

2. To remove the old certificate binding, run the following command:

   httpcfg.exe delete ssl -i 0.0.0.0:{CS SSL Port Number}

   The CS SSL Port Number parameter is the port number you entered during the Control Service installation. You can find this value in the ws_man.exe.config file, under ws_port value.

   The command result should not return any errors.

   The end of the message should be:

   ...completed with 0.

3. To bind the new certificate to the Control Service SSL port, run the following command:

   httpcfg.exe set ssl -i 0.0.0.0:{CS SSL Port Number} -h {New Certificate SslHash}

   The httpcfg.exe parameter is a standard utility for Windows Servers, and you can find it in the Control Service installation directory.

   You can find the New Certificate SSlHash parameter in the Certificate dialog, on the Details tab, under the Thumbprint value. Enter the Thumbprint value without spaces, in a single continuous string.

   The command result should not return any errors.

   The end of the message should be:

   ..completed with 0.

   The SSL certificate is now renewed.

# Unable to start listening on port/Ports

**"CS00073" "Unable to start listening on port %1 %%%"**

Occupied or closed by a firewall (Engine).

**"EM02012" "Unable to get port of Web Service on %1 %2"**

**Control Service Occupied or closed by a firewall - 8088**

**"EM02014" "Control Service has different port on %1 %2 and %3 %4 hosts"**

**"ER00609" "Failed to configure Web Access port."**

**Scenario IIS -**

Check whether another Engine or another application is already using this port. Change the port no. in one of them.

# Open ports required for remote installation and verification of the Engine

Title: List of open ports required for remote installation and verification of the Engine on replication hosts.

**Description**

This article provides the ports and the associated protocols in order to remotely install the Engine service on remote hosts through the scenario creation wizard or the remote installer.

**Solution**

The ports below need to be open in any firewalls between the Arcserve RHA Control Service server and the Engine servers.

- TCP Protocol requires ports 25000, 1025, 2666 & 2660

- UDP Protocol requires ports 135, 137, & 138

# Changing the Control Service Port

Arcserve RHA Control Service is designed by default to listen on port 8088. However, the default port can be changed in environments where this port is already being used by another application. The configuration file that is responsible to the Control Service port configuration is ws_man.exe.config. Therefore, if you want to change the Control Service port after you installed it, you need to change all the port values in the ws_man.exe.config file.

To change the default port of the Control Service from port 8088 to any other port

1. If there are running scenarios that are currently using the Control Service you want to change, stop them through Arcserve RHA Manager.

2. Log in to the host where the Control Service is running.

3. In the **Services** dialog, stop the Control Service.

4. Using Windows Explorer, browse to the Control Service installation directory, where the **ws_man.exe.config** file is located .

   **Notes**:

   - On a 32 bit host, the default installation directory is:"C:\Program Files\CA\Arcserve RHA\Manager"

   - On a 64 bit host, the default installation directory is: ""C:\Program Files (x86)\CA\Arcserve RHA\Manager"

5. Open the **ws_man.exe.config** file with a text editor.

6. Change the **value** of all '**\*_port** entries, and instead of the default port of 8088 enter the port no. you want to use.

7. Save and close the **ws_man.exe.config** file.

8. In the **Services** dialog, restart the Control Service.The Control Service now listens to the new port you defined for it.

# Changing the Engine Port

Arcserve RHA Engine is designed by default to listen on port 25000. However, the default port can be changed in environments where this port is already being used by another application. The configuration file that is responsible to the port configuration is ws_rep.cfg file. Therefore, if you want to change the Control Service port after you installed it, you need to change all the port values in the ws_rep.cfg file.

**To change the port number used by the Replication Engine**

1. If there are running scenarios that are currently using the Engines you want to change, stop them through Arcserve RHA Manager.

2. Log in to the Master host where the Engine is running. (Repeat for Replica)

3. In the **Services** dialog, stop the **Engine** service. (Stop the engine on both Master and Replica servers.)

4. Using Windows Explorer, browse to Engine installation directory, where the **ws_rep.cfg** file is located.

   **Note**: The default installation directory is: "C:\Program Files\CA\Arcserve RHA\Engine".

5. Open the **ws_rep.cfg** file with WordPad o another text editor.

   **Note**: We do not recommend using Notepad, due to its limited View options.

6.  WS_REP.CFG file using WORDPAD or a 3rd party text edition. (don't use NOTEPAD).

7.  Find the "# Port = 25000" section (one of the first lines) in the WS_REP.CFG file.

8.  Change the Port = 25000 to a new port number (e.g. Port = 25002) and remove # sign at the beginning of the line.

9.  Save the WS_REP.CFG (Remember: the configuration files on the Master and Replica servers must be the same. Make sure whatever changes that are made to the WS_REP.CFG file on the Master are made on the Replica.

10. Start the Engine Service on both the Master and Replica.

11. Open the Manager and highlight the scenario.

12. Click on the Active Server and then choose Properties. Under the connection section, you can see the port number and the value set to 25000. Change the port number to the new port number that is specified in the WS_REP.CFG file.

13. Perform the same steps that are performed in step 10 for the Replica server also.

14. If the Master and Replica servers have other scenario's running you will have to change the port number for those scenarios as well.

15. Restart the scenario.

**To change the default port of the Control Service from port 8088 to any other port**

1.  Log in to the host where the Control Service is running.

2.  In the **Services** dialog, stop the Control Service.

3.  Using Windows Explorer, browse to Control Service installation directory, where the **ws_man.exe.config** file is located .

    **Notes**:

    ◆ On a 32 bit host, the default location is:"C:\Program Files\CA\Arcserve RHA\Manager"

    ◆ On a 64 bit host, the default location is: "C:\Program Files (x86)\CA\Arcserve RHA\\Manager"

4.  Open the **ws_man.exe.config** file with a text editor.

5.  Change the **value** of all **\*_port** entries, and instead of the default port of 8088 enter the port no. you want to use.

6.  Save and close the **ws_man.exe.config** file.

In the **Services** dialog, restart the Control Service.The Control Service now listens to the new port you defined for it.

## Root Directories

"CV01361" " Invalid set of Root Directories"

"EM00568" "The host %1 is already used by running HA scenario '%2'. "

The %1 host is already used by a running HA scenario '%2'."

## Synchronization Failed with the VSS Error

**Symptom:**

I get the following error when I run synchronization:

WM04411 342 Warning 155.35.86.133 9/10/2012 5:17:00 PM Unable to create Shadow Copy; retrying...

**Solution:**

This error is because of specific hardware configuration. You can solve the problem by performing the following steps:

Open the ws_rep.cfg and modify the following parameters:

DirSnapshotWithVSS = True
MaxVSSRetryCount = 3

**Note:** Increase the MaxVSSRetryCount value if it is already 3.

## Unable to Create Shadow Copy Error

**Symptom:**

I get the following error while running the scenario:

Unable to Create Shadow Copy, error: 'The system or provider has insufficient storage space.'; retrying...

**Solution:**

1. Open the we_rep.cfg file and enable the SnapshotStorageVolume parameter and set its value to a volume with sufficient available space.

2. Restart the engine.

3. Run the scenario again.

# Unable to Create Shadow Copy Error for VHD Mount Point

**Symptom:**

I am unable to create shadow copy while synchronization when the VHD mount point coexists with other directories as the root directory.

**Solution:**

The problem is because there could be more than one root directory and the directories are in more than one volume. In this case, if any of the volumes are from a VHD file and the others are not, then the scenario cannot run as it cannot create a shadow copy set.

Perform the following steps to resolve the issue:

1. Open the ws_rep.cfg and modify the following parameters:

   DirSnapshotWithVSS = False

2. Run the scenario again.

# Index

## A

# D

# M

# N

# O

# P

# S

schedule  173, 271

**Switch Computer Name  325**

**Switchback**

Control Service roles  363

initiating  309

understanding  318

**Switchover**

active and standby host  318, 336

Active to Standby Redirection Script  325

add IP/Mask  325

automated  323

automatic or manual, defining  323

host name, defining  323

hosts, setting  324

Identify Network Traffic Direction Script  325

initiating  306

Is Alive check  331

Move IP  325

NetBIOS name for  325

perform automatically, define  323

properties  320, 323

Reboot after, setting  325

Redirect DNS  325

redirection methods, define  325

reverse replication scenario initiation, defining  323

Run Reverse Replication Scenario after, define  323

Standby to Active Redirection Script  325

Switch Computer Name  325

**Synchronization**

advanced scheduling  245

automatic  26

automatic, scheduling  242

Block  23

dialog  150

excluding dates from schedule  244

filter  25

hours  243

how it works  22

keep deleted files on Replica during  259

manual  150

method, selecting for run  144

methods  23

report see Synchronization report  255

# T

# U

# V

**Viewing**

events  186

events in separate window  187

incoming events using pop-up window  188

Manager screen options  50

reports  192

VSS Snapshots  393

**VSS Snapshots**

configuring creation  388

creation, setting up  389

Management window  393

managing  394

Max Storage Size per Volume  391

Preferred Number of Snapshots to Keep  391

properties  391

using  387

viewing  393

window  393

# W

**Wizard**

Restore data  295

Scenario Creation  66