

Arcserve® Replication and High Availability

Virtualized Server Environments Operation Guide for Windows

r16.5

arcserve®

Pre-release Document, only for reference

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

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

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

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

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

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

The manufacturer of this Documentation is Arcserve.

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

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

Arcserve Product References

This document references the following Arcserve products:

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

Contact Arcserve

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

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

With Arcserve Support:

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

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

Documentation Changes

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

- Updated to include user feedback, enhancements, corrections, and other minor changes to help improve the usability and understanding of the product or the documentation itself.
- Updated Configuring Master and Replica for Full System Scenarios section to include information for additional virtual platforms.
- Updated Protecting Full Systems section to include information for using a cloud destination.
- Updated VMware vCenter Server RHA Configuration section to include information for new vCenter Server 4.0 support.

Contents

Chapter 1: Introduction 7

About This Guide	7
Related Documentation	8
Log On Account Conditions	8
Register Arcserve RHA Licenses	8
Virtualization Considerations	9

Chapter 2: Protecting VMware vCenter Server Environments 11

VMware vCenter Server RHA Configuration	12
Configure the vCenter Server Master Server	14
Configure the vCenter Server Replica Server	14
vCenter Server Replication and High Availability	19
Create vCenter Server Replication Scenarios	19
Create vCenter Server HA Scenarios	21
Set Scenario Properties	22
The Data Recovery Process	25
Redirection Methods	41
How Redirection Works	41
DNS Redirection	42
Move IP Redirection	42
Switch Computer Name Redirection	46
Scripts Redirection	47
Switching Over and Switching Back	48
How Switchover and Switchback Work	49
Initiate Switchover	50
Initiate Switchback	52
Switchover Considerations	54
Run a Scenario	54
Stop a Scenario	56
View a Report	56

Chapter 3: Protecting Hyper-V Environments 59

Hyper-V Server Configuration Requirements	59
Hyper-V HA Configuration	60
Hyper-V VM Auto-Discovery	61
Configure Hyper-V for High Availability	61

Considerations for Hyper-V scenarios across WAN	67
Hyper-V Replication and High Availability.....	68
Create a New Hyper-V Replication Scenario	68
Hyper-V Replication Properties.....	69
Hyper-V HA Properties	70
Create a New Hyper-V High Availability Scenario	73
Switching Over and Switching Back.....	75
How Switchover and Switchback Work.....	76
Initiate Switchover	77
Initiate Switchback	79
Switchover Considerations.....	81
Run a Scenario from Outside the Wizard	82
Stop a Scenario.....	83
View a Report.....	83
The Data Recovery Process	84
Setting Bookmarks	85
How to Restore Data on Hyper-V Machines	86
Additional Information and Tips	88
Spool Directory Settings.....	90
Manage Services	91
Chapter 4: Protecting Full Systems	93
How Full System Scenarios Work	93
Unified Extensible Firmware Interface (UEFI) Support for Full System Scenarios	96
Index	97

Chapter 1: Introduction

Arcserve Replication and High Availability (Arcserve RHA) is a solution based on asynchronous real-time replication and automated application switchover and switchback to provide cost-effective business continuity for virtualized environments on Windows servers.

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

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

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

This section contains the following topics:

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

[Related Documentation](#) (see page 8)

[Log On Account Conditions](#) (see page 8)

[Register Arcserve RHA Licenses](#) (see page 8)

[Virtualization Considerations](#) (see page 9)

About This Guide

This document describes how to implement a Arcserve Replication and High Availability solution for virtualized environments. Please review each procedure before you begin. It is essential that you have the appropriate resources and permissions to carry out each task.

The Guide is organized into the following main sections:

- **Microsoft Windows Hyper-V** -- Full server replication and switchover for each individual guest machine
- **VMware vCenter Server** -- vCenter Management Console Replication and switchover
- **Full System High Availability (HA)** -- Transfers an entire physical system to a virtual machine hosted by a Hyper-V server

Related Documentation

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

Log On Account Conditions

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

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

Note: In MS SQL server 2012, Local System (NT AUTHORITY\SYSTEM) is not automatically provisioned in the sysadmin server role. See Microsoft documents on how to provision the sysadmin server role to the NT AUTHORITY\SYSTEM account. Alternatively, use Administrator account to install and log in to the engine service.

- When your SQL servers are in a workgroup, enable "sa" account on Master and Replica server before you run the scenario.

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.

Virtualization Considerations

In virtual machine environments, all disks are added to the first controller after switchover, until it is full (16 disks). When it is full, Arcserve RHA creates another SCSI controller. This means the switched over VM could have a different deployment than its master.

Chapter 2: Protecting VMware vCenter Server Environments

This section contains the following topics:

[VMware vCenter Server RHA Configuration](#) (see page 12)

[vCenter Server Replication and High Availability](#) (see page 19)

[Redirection Methods](#) (see page 41)

[How Switchover and Switchback Work](#) (see page 49)

VMware vCenter Server RHA Configuration

vCenter Server has several components (Database Server, License Server and Web Access Server) that can be deployed to a single machine (local) or in a distributed manner.

- Install the same VMware vCenter Server components on the Replica so they match the components on the Master.
- Install the Arcserve RHA Engine on every vCenter Server.
- (In the Active Directory environment) Both Master and Replica servers should reside in the same Active Directory forest and be members of the same domain or trusted domains.

The following describes the two main methods of configuring the vCenter Server:

- **Single machine (local) deployment**--the Database Server, vCenter Server, License Server and Web Access Server are all installed on the same machine. If you install these components on the Master server, you must also install them on the Replica server.
- **Distributed deployment**--vCenter Server and Web Access Server are installed on one machine, while the Database Server, the License Server, or both are installed on another. If the vCenter Server machine goes down, the database can still operate. If you use distributed deployment, you must ensure the settings on both the Master and Replica servers both point to the same distributed machine(s). For distributed deployment, create a SQL or Oracle HA scenario to protect the separate database. See the appropriate Operation Guide for more information.

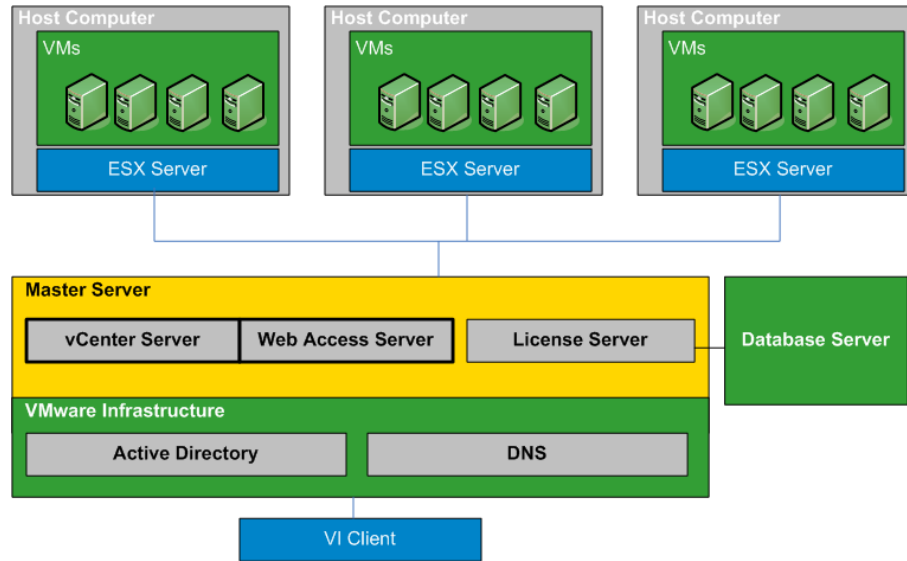
Important! Regardless of the method you choose, the vCenter Server and the Web Access Server must be installed on the same machine.

If License Server deployment is distributed, perform one of the following:

- If License Server (A) is installed on the Master, install another instance of License Server (B) on the Replica and specify the B instance as the License Server for the vCenter Server running on the Replica.
- If License Server (A) is installed on a machine other than the Master server, specify the first instance (A) as the License Server for the vCenter Server running on the Replica.

Note: If you are running VMware vCenter 4, there is no longer a separate License Server to install.

In the following diagram, VMware vCenter Server is configured using the distributed deployment method. The vCenter Server system that houses the vCenter Server and Web Access Server is the Master Server. You must install the Arcserve RHA Engine on each vCenter Server in your environment:



Configure the vCenter Server Master Server

Consider the following points when configuring your vCenter Server environment:

- Each VMware ESX server locally stores the IP address of the vCenter Server managing it. This address must be changed to the standby server in the event of a switchover. We recommend using the vCenter Server Managed IP Address setting to automate switchover.
- Each VMware ESX server can be managed by only one vCenter Server. Connecting an ESX server to more than one vCenter Server automatically disconnects it from its original server.

When you configure Arcserve RHA on the vCenter Server Master Server, enter the same IP address you entered in the Managed IP Address field as the IP Mask field on the High Availability Properties screen.

If you do not use the Managed IP Address setting, you must manually reconnect all ESX servers to the standby vCenter Server after switchover.

To configure the vCenter Server Master server

1. Add another IP address to the Master server network interface card (NIC). For more information, refer to the topic, Add IP on the Master Server.
2. From the vCenter Server Runtime Settings screen, set the Managed IP address field to the same IP address you added in Step 1, not the Master server's IP address. For more information, refer to the VMware vCenter Server documentation.
3. Run the HostReconnect script to reconnect all ESX servers managed by the vCenter Server. Download this script from the VMware website.

Important! When you create HA scenarios, we recommend that you enable the Move IP redirection method on the Switchover Properties screen. If you enter additional IP addresses to the Master Server NIC, Move IP is enabled by default.

Configure the vCenter Server Replica Server

To set up VMware vCenter server for replication, configure a Replica server as follows:

- Install vCenter components identically on Master and Replica servers.
- If you are using distributed database deployment, configure vCenter server on the Replica to connect to the same Database Server configured for Master. If you are using local database deployment, configure vCenter server on the Replica to connect to the Database Server configured on the Replica.

- Configure vCenter server on Replica to connect to License Server: If License Server is installed locally to the Master server, you must install another instance of License Server on the Replica and specify this instance in vCenter on the Replica. If the License Server is installed remotely to the Master, specify that instance in vCenter on the Replica.
- Install the Engine on every vCenter server.

Note: For distributed database deployment, protect the database using Arcserve RHA for SQL Server or Oracle, as appropriate. Refer to the Operation Guide for more information. For distributed License Server deployment, the License Server cannot be protected in Arcserve RHA scenarios.

To set up VMware vCenter Server for high availability, configure a Replica server as follows:

- Install vCenter Server components identically on Master and Replica servers. Ensure the Replica has the same database type as the Master. You should also use the same folder structure on both servers.
- Configure the Database Server on the Replica according to the vCenter Server configuration method you used. For more information, see [Configure the vCenter Server Database Server](#). (see page 16)
- Configure the License Server on the Replica according to the vCenter Server configuration method you used. For more information, see [Configure the vCenter Server License Server](#). (see page 18)

Note: To protect the distributed (remote) database servers, create separate Arcserve RHA scenarios for SQL or Oracle, as appropriate. However, if the License Server is deployed separately, Arcserve RHA cannot protect it.

Configure the vCenter Server Database Server on the Replica

In the event of a failure, access to the Database Server is essential to a VMware vCenter Server high availability scenario.

Note: Specify the server name explicitly to prevent database auto-discovery problems.

To configure the vCenter Server Database Server on the Replica server

Do one of the following, depending on how your environment is configured:

- If you used single machine (local) deployment, configure the ODBC settings as follows:
 1. Install the Database Server on the Replica. Explicitly specify this instance on the Replica so that the Replica uses its local Database Server.
 2. Ensure the Database Server Instance Name is the same as the one specified on the Master.
 3. Ensure the Database Name is the same as the one specified on the Master.
 4. Ensure the Database Server installation path and database files path are the same as those specified on the Master.

Or:

- If you used distributed deployment, configure the ODBC settings as follows:
 1. Do not install the Database Server on the Replica. Instead, specify the same remote Database Server that is specified on the Master.
 2. Specify the same data source name (DSN) as on the Master. To do this, select **Use an existing database server** from the VMware vCenter Server Deployment Options - Step 1 screen and then enter the name of an already-configured DSN. Click No when asked if you wish to re-initialize the database and start over with a blank configuration.

Supported Databases VMware vCenter Server

Arcserve RHA supports the following databases when used in a VMware vCenter Server environment:

- Microsoft SQL Server 2005, 2008
- Microsoft SQL Server Express
- Oracle 10g, 11g

If the vCenter Server Database Server is deployed separately from the vCenter Server, you must create additional scenarios using the appropriate database application type to protect it. Database Servers are protected in vCenter Server scenarios only when deployed on the same machine hosting the vCenter Server.

Note: If you wish to protect MSDE and Microsoft SQL Server 2005 Express, visit the Microsoft technical support website and download the knowledge document, TEC445313.

Configure the vCenter Server License Server on the Replica

In the event of a failure, access to the License Server is also essential to a vCenter Server high availability scenario. Perform one of the following procedures if you are not running vCenter Server 4. There is no longer a separate license server in vCenter Server 4 deployments.

To configure the vCenter Server License Server on the Replica server:

Do one of the following:

- If you used single machine (local) deployment:
 1. Install the License Server on the Replica.
 2. Replica uses its local License Server. From the VMware vCenter Server Deployment Options-Step 2 screen, choose Use an Existing License Server.
 3. Specify the Replica License Server and port number in the text field.

Or:

- If you used distributed deployment:
 1. You do not need to install the License Server on the Replica.
 2. From the VMware vCenter Server Deployment Options-Step 2 screen, choose Use an Existing License Server
 3. Specify the same remote License Server that is specified on the Master in the text field.

Next, configure the SSL Certificate

1. Copy the SSL Certificate from the Master server to the Replica server.
2. Use the following command to input a user name and password on the Replica server to reset encrypted passwords in the Registry.

```
vxdp -p
```

The Database server can now be accessed by vCenter Server on the Replica server.

vCenter Server Replication and High Availability

Create vCenter Server Replication Scenarios

VMware vCenter Server application data is file-based and is auto-discovered during scenario creation.

To create a new vCenter replication scenario

1. Open the Manager and choose Scenario, New or click the New Scenario button to launch the Scenario Creation wizard.
The Welcome dialog opens.
2. Choose Create a New Scenario, select a Group Name and click Next.
The Select Server and Product Type dialog opens.
3. Select vCenter, Replication and Disaster Recovery Scenario (DR) and click Next.
Note: If you wish to specify Tasks on Replica, see the Arcserve RHA Administration Guide for more information.
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, enable the Verify Engine on Hosts option and then click Next.
Wait for Engine Verification to complete.
5. Click Install to upgrade the Engine service on one or both servers, if necessary, and then click Next.
The Master Configuration dialog opens.
6. If your vCenter database is SQL Server, discovery results are shown in the DB tree. If your vCenter database is Oracle, you are prompted for Database Connection credentials. Provide the required information and click OK.
Clear or check components for disaster recovery, as desired, and then click Next.
The Scenario Properties dialog opens.
7. Set the desired properties as described in the Arcserve RHA Administration Guide and click Next.
The Master and Replica Properties dialog opens.
8. Set the desired properties and click Next.
Wait for Scenario Verification to complete.
9. Resolve any warnings or errors reported at Verification and then click Next.
The Scenario Run dialog opens.

10. Click Run Now to initiate replication and activate the scenario or click Finish to run the scenario later.

Create vCenter Server HA Scenarios

During discovery, system databases are automatically protected. However, if the Database Server used by vCenter is a remote machine, this scenario cannot protect it. You must create a specific database scenario to protect a remote Database Server.

To create a new vCenter Server HA scenario

1. Open the Manager and click Scenario, New or click the New Scenario button.
The Welcome dialog opens.
2. Choose Create a New Scenario and select a Scenario Group from the list. Click Next.
The Select Server and Product Type dialog opens.
3. Choose VMware vCenter Server, High Availability Scenario (HA) and 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 the Master and Replica servers, enable the Verify Engine on Hosts option, and then click Next.

Wait while Engine Verification completes. If prompted for logon credentials, enter them and click OK.

5. If necessary, click Install to upgrade the Engine service on one or both Servers. If you are prompted for logon credentials, enter them and click OK. When installation completes, verify again and then click Next.

Note: If you are using a local Oracle instance as the vCenter database, you are now prompted for logon credentials. Enter the Oracle DBA name and password now and click OK to start auto-discovery.

The Database for Replication dialog opens and displays the auto-discovered results for the host you specified as the Master. For SQL Server, all databases used by vCenter are replicated by default.

6. Clear checkmarks next to the databases you do not want to replicate, if desired, and click Next.

The Replica Configuration dialog opens. Arcserve RHA automatically compares Master and Replica configurations, ensuring they are identical.

7. Resolve errors, if any are displayed and then click Next.

The Scenario Properties dialog opens.

8. Configure additional properties, if desired, and click Next. For more information, see [Scenario Properties](#) (see page 22) or the Administration Guide.

The Master and Replica Properties dialog opens.

9. Make changes, if desired, and click Next. For more information, see Scenario Properties or the Administration Guide.

Wait while the Switchover Properties dialog retrieves information.

10. Configure the desired redirection properties, and click Next. For vCenter HA scenarios, Move IP is automatically set to On. For more information, see [Switching Over and Switching Back](#). (see page 48)

The Switchover and Reverse Replication Initiation dialog opens.


11. Choose automatic or manual switchover, and automatic or manual reverse replication, as desired, and click Next. You should not set both of these options to Automatic. For more information, see Scenario Properties or the Administration Guide.
12. Wait while Scenario Verification completes. Problems are reported in the dialog. You must resolve errors and should resolve any warnings before running the scenario. When scenario verification is successful, click Next.
13. Choose Run Now to start synchronization or Finish to save the scenario and run it later. For more information, see Run the Scenario from Outside the Wizard.

Set Scenario Properties

You can change a scenario configured using the Wizard or configure additional settings, or modify the scenario, using the Properties pane.

The Properties pane and its tabs are context-sensitive and change whenever you select a different node from a scenario folder. You must stop a scenario before configuring its properties. Certain values cannot be modified once set; they are noted. For full details on configuring scenario properties and their descriptions, see the *Arcserve RHA Administration Guide*.

Properties are organized into tabs on the Arcserve RHA Manager Framework pane. The tabs displayed are based upon server type, Arcserve RHA solution, and scenario status. Select the scenario for which you want to change properties, and then select the appropriate tab.



Settings on the Root Directories tab

Do the following:

1. Select a Master Server from the Scenario Pane. Double-click its Directories folder to add or remove Master Root Directories. Select or clear checkboxes next to folders, as desired, to include or exclude them. You may also edit directory names.
2. Select a Replica Server from the Scenario Pane. For each Master Root directory, you must specify a Replica Root directory. Double-click the Directories folder for the Replica server. Select or clear checkboxes next to folders, as desired, to hold the corresponding Master directory.

Settings on the Properties Tab

Scenario Properties

These settings establish default behavior for the entire scenario.

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

Master and Replica Properties

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

- Host connection properties -- Enter the IP address, Port number and Fully Qualified Name of the Master and Replica.
- Replication properties -- These properties differ for Master and Replica. See the *Arcserve RHA Administration Guide* for more information.
- Spool properties -- Set the size, minimum disk free size and directory path. See [Spool Directory Settings](#) (see page 90) for more information.
- Event notification properties -- specify a script to run, select email notification, or write to event log.
- Report properties -- select synchronization or replication reports, specify distribution or script execution.
- (Replica) Scheduled Tasks -- set or suspend tasks, including Replica Integrity Testing for Assured Recovery. For more details, see the *Arcserve RHA Administration Guide*.
- (Replica) Recovery properties -- set delay, data rewind properties, or scheduled task for replica.

Settings on the HA Properties Tab

These settings control how switchover and switchback are performed.

- Switchover properties -- select automatic or manual switchover, provide switchover hostname, and reverse replication settings.
- Hosts properties -- specify the Master and Replica Fully Qualified Name.
- Network Traffic Redirection properties -- select Move IP, Redirect DNS, Switch Computer Name or User-defined scripts.
- Is Alive properties -- set the heartbeat frequency and check method.
- DB Management properties (does not apply to File Server scenarios) -- instructs Arcserve RHA to manage shares or services on a database server.
- Action upon Success properties -- defines custom scripts and arguments for use.

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

The following procedure is the same for all server types.

To recover all lost data from a Replica

1. On the Manager, select the desired scenario from the Scenario pane and stop it.
2. (For database applications only) Stop the database services on the Master host.
3. On the Manager, select the Replica host from the scenario folder. If multiple Replicas participate in a scenario, select the one from which you wish to recover data. The Restore Data option is now enabled.
4. From the Tools menu, select Restore Data or click the Restore data from the Standard toolbar.

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 opens, asking you to enter logon account details for the selected Replica.

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

Note: If the Data Rewind property is set to On, another Restore Data dialog will appear. In this case, select the first option - **Replace all data on Master with the data on Replica**. This option simply restores data without a rewind.

5. Click **Next**. The **Synchronization Method** page appears.
6. Make sure that the appropriate Synchronization method is selected. For more details, see the *Arcserve RHA Administration Guide*. Click **Finish**.

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

Synchronization finished.

Note: If an error occurred during the temporary Recovery scenario run, the Recovery scenario may stop and remain in the Scenario pane. In this case, you should remove it by right-clicking it and selecting the **Remove** option from the pop-up menu. After the Recovery scenario is removed, the original scenario re-appears in the Scenario pane. Then, you can restart the original scenario, and repeat the recovery process if necessary.

7. By default, once a data recovery occurs a Synchronization Report is generated.
Now, the replication process can restart following the original scenario.

Recover Active Server

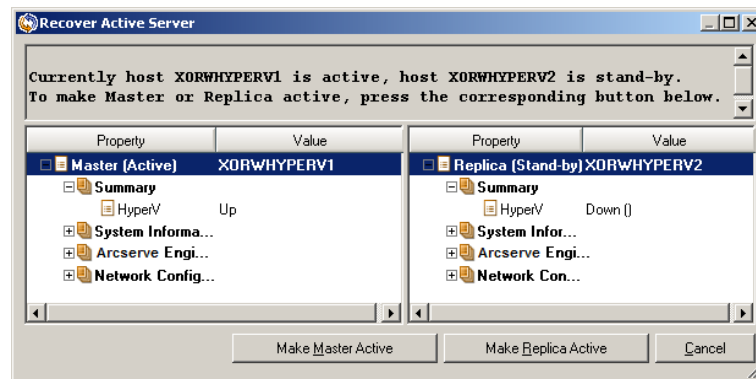
In certain circumstances, it may be necessary to forcibly make the Master or Replica server the active server without completing the data synchronization process.

For example, if switchover occurred but no data was changed on the Replica server. In this case you may even have newer data on the Master server making it undesirable to synchronize data from the Replica to the Master server. Arcserve RHA allows Recover Active Server process, to use this option. You must ensure that the scenario is stopped, and select *Recover Active Server* from the Tools menu.

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

If the Recover Active Server method does not solve the problem, you can manually recover a server. For more information, refer to the section, [Recovering Servers](#) (see page 31).

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



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

Setting Bookmarks

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


Notes:

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

To set a bookmark

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

The Rewind Bookmark dialog opens.



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

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

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

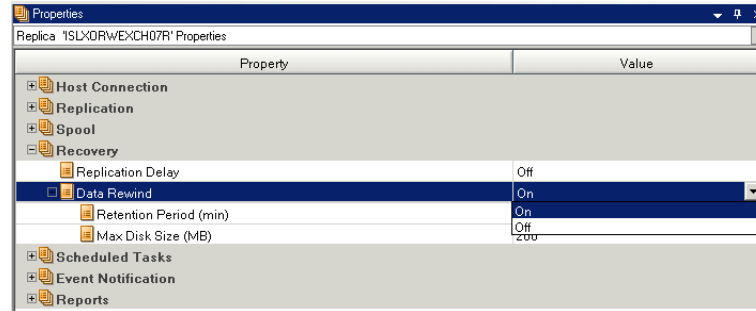
The bookmark is set.

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

Data Rewind

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

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



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


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


Note: The automatic registration of the rewind points starts only after the synchronization process is completed, and the message All modifications during synchronization period are replicated appears on the Event pane. Similarly, you cannot manually set bookmarks during synchronization. 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. Select the scenario that you want to stop from the Scenario pane on the Manager and stop it.
2. [For database applications only] Stop the database services on the Master host.
3. Select the Replica host from the scenario folder:

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



4. Select Restore Data from the Tools menu or click the Restore Data  button. If you are prompted for user credentials, enter the appropriate information and click OK.

The Recovery Method page of the Restore Data Wizard opens.

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 check box is enabled only if you activated this option before starting the scenario. If the check box is enabled, you can select it to include the synchronized Registry Keys in the recovery process.

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

6. Click Next.

The Rewind Point Selection page is displayed.

7. Wait until the Select Rewind Point button is enabled, and click it to view the existing rewind points.

The Select Rewind Point dialog opens.

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

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

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, which now displays information about the rewind point you selected.

9. Click Next.

The Synchronization Method page is displayed.

10. Select the Block Synchronization method and click Finish.

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

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

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

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

How to Restore Data on vCenter Machines

Using rewind points, vCenter Server has its own view on the Select Rewind Point dialog. Click the vCenter option to enable this view and then select a rewind point from the desired machine.

Recovering Servers

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

- Perform the Recover Active Server procedure. For more information, see Recover Active Server.
- If the Recover Active Server procedure does not resolve the issue, try one or more of the following manual tasks appropriate to the redirection method you use:
 - If IP Redirection is used, manually remove the IP. You cannot use this method for scenarios that do not support Move IP redirection (Hyper-V HA, CS HA). For more information, see [Manually Recover a Failed Server-Move IP Address](#). (see page 32)
 - If Switch Computer Name Redirection is used, manually switch the names. You cannot use this method for scenarios that do not support Switch Computer Name Redirection (Hyper-V HA, Exchange HA, vCenter HA if local Oracle is used). For more information, see [Manually Recover a Failed Server-Switch Computer Name](#) (see page 32).
 - If both IP and Switch Computer Name Redirection methods are used, manually remove the IP and switch the computer names. You cannot use this method for scenarios that do not support Move IP and Switch Computer Name redirection (Exchange, CS HA). For more information, refer to the topic, [Manually Recover a Failed Server-IP and Switch Computer Name](#) (see page 33).

Manually Recover a Failed Server-Move IP Address

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

To recover a failed server using Move IP Address redirection method

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

Manually Recover a Failed Server-Switch Computer Name

Important! When Oracle is installed locally and used by vCenter Server, the Switch Computer Name redirection method is not supported.

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

1. Boot the Master server without a network connection, to avoid duplicate network names.
2. Rename the server to <NewServerName>-RHA and move it to a temporary workgroup. For example, if the server is called "Server1", rename it to "Server1-RHA". You will be required to reboot this machine. After the reboot completes, the following error appears: "At least one Service could not be started." Ignore this, it is normal under these circumstances because the Engine usually runs in a domain account.
3. Connect to the network.
4. Rejoin the domain, ensuring that you use the -RHA name assigned in step 2.
5. Reboot the computer.
6. If it is not already running, start the scenario from the Manager. (If automatic reverse replication was set to On, the scenario runs in backward mode so that the Replica server is now active and the Master server is now standby.)
7. Wait for synchronization to complete. Perform a manual switchover to make the Master server active. It is recommended that you do so outside of normal business hours.

Manually Recover Failed Server-IP and Switch Computer Name

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

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

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

Troubleshooting vCenter Server Scenarios

The following information is provided to help you resolve any errors and warnings.

EM00589 License Servers configured on vCenter are not consistent

Reason:

License Server deployment on the Master and Replica is not the same. One is locally installed and the other is remotely installed. These should be the same on Master and Replica.

Action:

Reconfigure the settings for the License Server on the Replica.

EM00590 Databases configured on vCenter are not consistent

Reason:

Either the deployment or the database type of the database server on the Master differs from the Replica and must be the same.

Action:

Reconfigure database settings on the Replica.

EM00591 vCenter versions configured are not consistent

Reason:

The vCenter version on the Master differs from that on the Replica and must be identical.

Action:

Re-install vCenter on the Replica.

EM00592 Distributed databases configured on vCenter are not consistent

Reason:

The Database Server hostname or instance name on the Master differs from that on the Replica and must be the same.

Action:

Reconfigure the database settings on the Replica.

EM00594 Database usernames configured for vCenter are not identical

Reason:

The account for vCenter used to access the Database Server on the Master differs from that on the Replica and should be identical.

Action:

Reconfigure the ODBC settings on the Replica.

EM00596 Oracle server name configured on replica for vCenter is not localhost

Reason:

This is a configuration error.

Action:

Change the ODBC setting using "localhost" as the hostname on the Replica.

EM00597 SQL server name configured on replica for vCenter is not localhost

Reason:

This is a configuration error.

Action:

Change the ODBC setting using "localhost" as the hostname on the Replica.

EM00598 The vCenter database names configured are not consistent

Reason:

This is a configuration problem.

Action:

Reconfigure the database settings on the Replica to match those on the Master.

EM00599 The vCenter database on master <IP Address> is distributed deployed and AR isn't supported

Reason:

If the vCenter Server database is on a remote machine, AR is not supported.

Action:

Restart Arcserve RHA. AR options should be dimmed when Arcserve RHA detects remote database deployment.

ER00603 Unknown parameter. vCenter configurations can't be compared

Reason:

This is due to an internal program error.

Action:

Retry.

ER00604 vCenter HA information is not initialized

Reason:

The Engine on the Master or Replica is having problems.

Action:

Verify the Engine is working and retry.

EM00590 Databases configured on vCenter are not consistent

Reason:

The Database server type on the Master differs from that on the Replica and must be the same.

Action:

Reconfigure the database server on the Replica.

ER00605 Failed to configure DB Server

Reason:

Something is wrong with the registry setting on the Replica:
HKEY_LOCAL_MACHINE_SOFTWARE\VMware, Inc.\VMware vCenter\DB

Action:

Confirm ODBC settings for vCenter are correct and check the registry key. If it does not exist, re-install vCenter on the Replica.

ER00606 Failed to configure License Server

Reason:

Something is wrong with the registry setting on the Replica:
HKEY_LOCAL_MACHINE_SOFTWARE\VMware, Inc.\VMware vCenter\vCenter

Action:

Check the registry key. If it does not exist, re-install vCenter on the Replica.

ER00607 Failed to configure Managed IP

Reason:

Something is wrong with the registry setting on the Replica:
HKEY_LOCAL_MACHINE_SOFTWARE\VMware, Inc.\VMware vCenter\vCenter

Action:

Check the registry key. If it does not exist, re-install vCenter on the Replica.

ER00608 Failed to configure DB password

Reason:

Something is wrong with the registry value "3" on the Replica:
HKEY_LOCAL_MACHINE_SOFTWARE\VMware, Inc.\VMware vCenter\DB

Action:

Check the registry key. If it does not exist, re-install vCenter on the Replica.

ER00609 Failed to configure Web Access Port

Reason:

Something is wrong with the registry value "WebCenterPort" on the Replica:
HKEY_LOCAL_MACHINE_SOFTWARE\VMWare, Inc.\VMware vCenter

Action:

Check the registry key. If it does not exist, reinstall vCenter on the Replica.

WM00529 Distributed License Servers configured on vCenter are not consistent

Reason:

The registry value "License Path" of key HKEY_LOCAL_MACHINE_SOFTWARE\VMware, Inc.\VMware vCenter on Replica should be the same as on Master.

Action:

Reconfigure the license server setting on the Replica.

WM00531 License Servers configured on vCenter are not consistent

Reason:

License Server deployment on the Master and Replica is not the same. One is locally installed and the other is remotely installed. These should be the same on Master and Replica.

Action:

Reconfigure the settings of the License Server on the Replica.

WM00532 Databases configured on vCenter are not consistent

Reason:

Either the deployment or the database type of the database server on the Master differs from the Replica and must be the same.

Action:

Reconfigure database settings on the Replica.

WM00533 vCenter versions configured are not consistent

Reason:

The vCenter version on the Master differs from that on the Replica and must be identical.

Action:

Re-install vCenter on the Replica.

WM00534 Distributed databases configured on vCenter are not identical.

Reason:

The Database Server hostname or instance name on the Master differs from that on the Replica and must be the same.

Action:

Reconfigure the database settings on the Replica.

WM00535 Unable to receive vCenter information from <IP Address>

Reason:

This is an internal program error, engine disconnect or timeout.

Action:

Wait and retry the request later.

WM00536 Database usernames configured for vCenter are not identical

Reason:

The account for vCenter used to access the Database Server on the Master differs from that on the Replica and should be identical.

Action:

Reconfigure the ODBC settings on the Replica.

WM00537 WebCenter ports configured on vCenter are not identical

Reason:

The WebCenter ports on the Master differ from those set on the Replica and should be identical.

Action:

Re-install vCenter on the Replica and ensure the WebCenter ports are the same as those on the Master.

WM00538 The vCenter Managed IP <IP Address> isn't set in Move IP list

Reason:

You have set a vCenter Managed IP but not added it to the Move IP properties located in the Switchover properties list.

Action:

Add the managed IP address to the Move IP list when setting switchover properties.

WM00540 SQL server name configured on replica for vCenter is not localhost

Reason:

This is a configuration error.

Action:

Change the ODBC setting using "localhost" as the hostname on the Replica.

WM00541 License server name configured on replica for vCenter is not localhost

Reason:

This is a configuration error.

Action:

Change the value "License Path" to "xxxx@localhost" style on Replica.

WM00542 License server ports configured for vCenter are not consistent

Reason:

This is a configuration error.

Action:

Re-install vCenter to reconfigure the license server on the Replica.

WM00543 License files folders configured are not consistent

Reason:

This is a configuration problem.

Action:

Reinstall vCenter to specify the correct folder for License files.

WM00544 The vCenter database names configured are not consistent

Reason:

This is a configuration problem.

Action:

Reconfigure the database settings on the Replica to match those on the Master.

WM00588 Distributed License Servers configured on vCenter are not consistent

The registry value "License Path" of key HKEY_LOCAL_MACHINE\SOFTWARE\VMware, Inc.\VMware vCenter on Replica should be the same as on Master.

Action:

Reconfigure the license server setting on the Replica.

Troubleshooting Oracle Databases

Oracle database fails to start after switchover

Symptom:

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

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

Solution:

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

1. Shut down

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

2. Start again

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

Redirection Methods

How Redirection Works

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

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

DNS Redirection

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

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

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

Move IP Redirection

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

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

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

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

Add IP on the Master Server

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

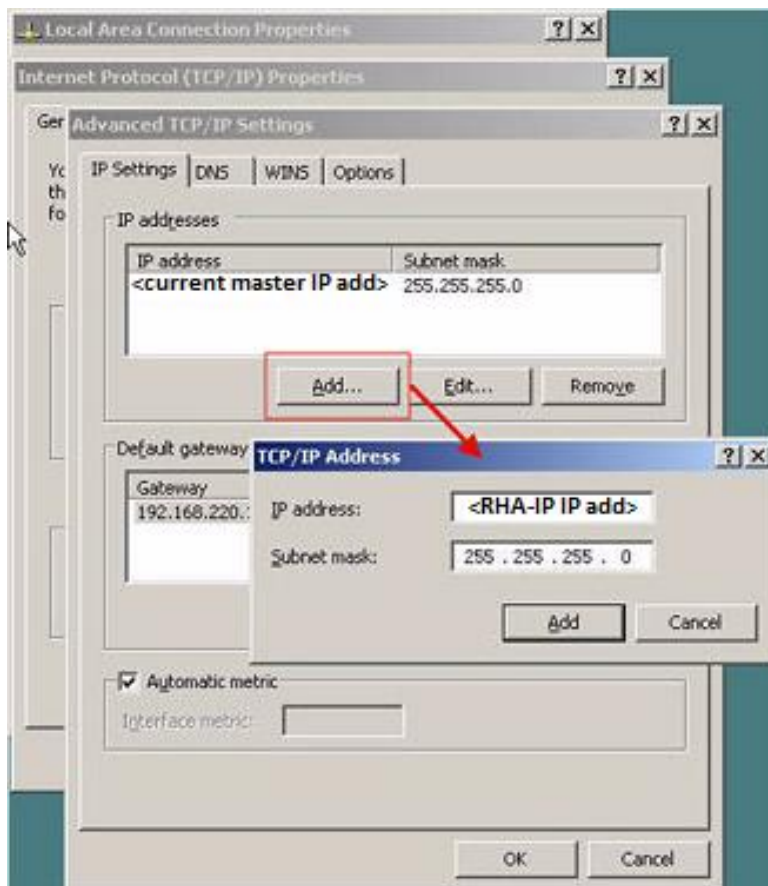
Important! Perform the following only if you are using the Move IP redirection method.

To add IP Address to Master Server

1. Open the Control Panel and select Network Connections.
2. Right-click Local Area Network and select Properties.

3. Click Internet Protocol (TCP/IP) and then click the Properties button.
4. Click Advanced.
5. Click Add and enter an additional IP address (Arcserve-IP).

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



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

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

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


The procedures for both ways follow.

Add Arcserve-IP to Existing Scenarios

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

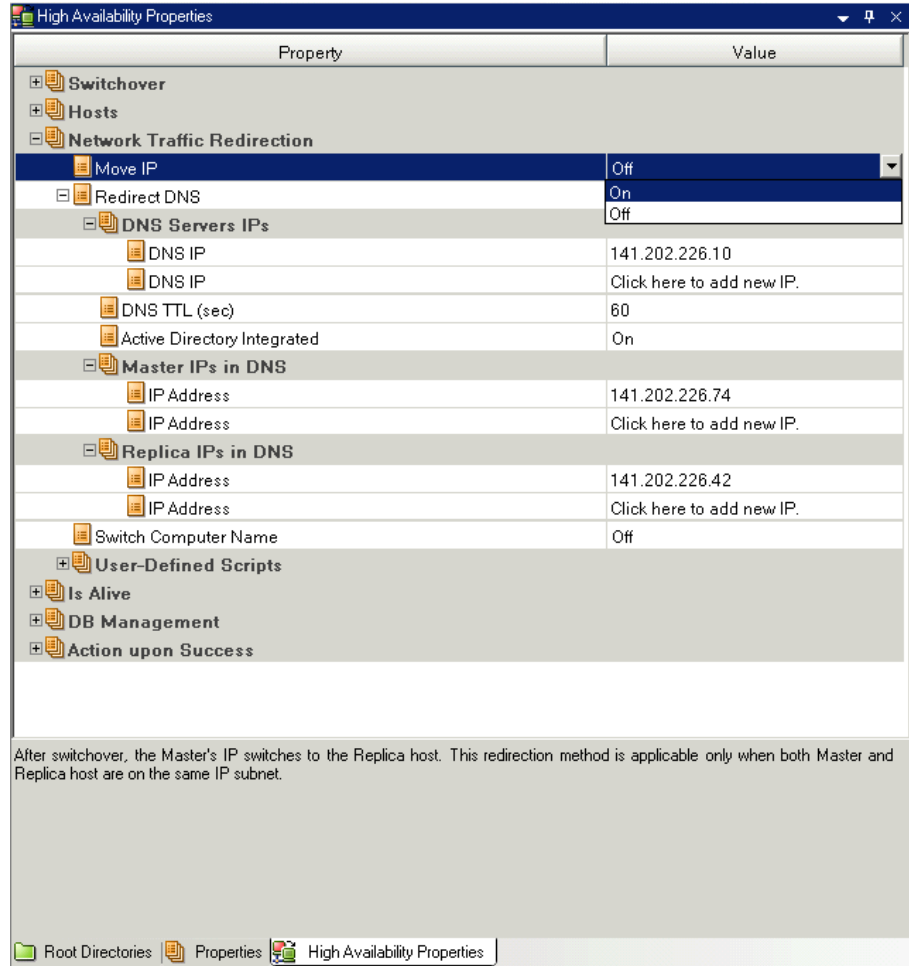
To add the Arcserve-IP to existing scenarios:

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



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


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 multiple production IP addresses by selecting Click here to add new IP/Mask.



Add Arcserve-IP to New Scenarios

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

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



Switch Computer Name Redirection

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

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

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

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

Automatic Redirection Using Switch Computer Name

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

Scripts Redirection

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

Because of the different ways in which a vCenter Server environment can be configured, redirection must ensure that the vCenter Server, Web Access, Database, and License Servers can all be accessed in the event of a failure. We recommend that you use the vCenter Managed IP feature. If the IP address of a vCenter Server system changes, the managed hosts connected to it are automatically reconnected if you use this feature.

If you are using this redirection method	Under this condition	Perform this task
Switch Computer Name	If you are using SQL Server as the database and the Database Server is on the Master...	...no additional action is needed. Auto Configure changes the ODBC setting on the Replica to the database server local to the Replica.
	If you are using Oracle as the database and the Database Server is on the Master...	...no additional action is needed.

If you are using this redirection method	Under this condition	Perform this task
	If the Database Server (SQL or Oracle) is on a remote machine...	...no additional action is needed.
	If the License Server is on the Master...	...change the License Path on the Replica to point to the local Replica License Server in registry: HKEY_LOCAL_MACHINE\SOFTWARE\VMware, Inc.\VMware vCenter Server key value: "LicensePath"="xxxxx@localhost"
	If the License Server is on a remote machine...	...no additional action is needed.
Move IP Address	If you are using the IP address to access vCenter Server...	...no additional action is needed. Auto Configure changes the path to the License Server in the registry.
DNS	If you are using the hostname to access the vCenter Server...	...enable the DNS redirection mode.

Switching Over and Switching Back

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

How Switchover and Switchback Work

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

- Ping--a request sent to the Master to verify that the Master is up and responding
- Database check--a request that verifies the appropriate services are running and all databases are mounted
- User-defined check--a custom request you can tailor to monitor specific applications

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

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

- If you selected the Initiate Switchover manually option from the Switchover and Reverse Replication Initiation page, perform a manual switchover. For more information, refer to the topic, [Initiate Switchover](#) (see page 50).
- If you selected the Initiate Switchover automatically option, you can still perform a manual switchover, even if the Master is alive. You can initiate switchover when you want to test your system, or you want to use the Replica server to continue the application service while some form of maintenance is performed on the Master server. Triggered (automatic) switchover is in all ways identical to manual switchover performed by the administrator, except it is triggered by a resource failure on the master server rather than by an administrator manually initiating the switchover by clicking the Perform Switchover button. The timeout parameters are configurable and are more extensively covered in the *Arcserve RHA Administration Guide*.

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

- If you selected the Initiate Reverse Replication automatically option from the Switchover and Reverse Replication Initiation page, replication in the reverse direction (from Replica to Master) automatically begins after a switchover has finished successfully.
- If you selected the Initiate Reverse Replication manually option, you must resynchronize data from Replica to Master, even after testing a clean switchover without a Master failure.

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

Initiate Switchover

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

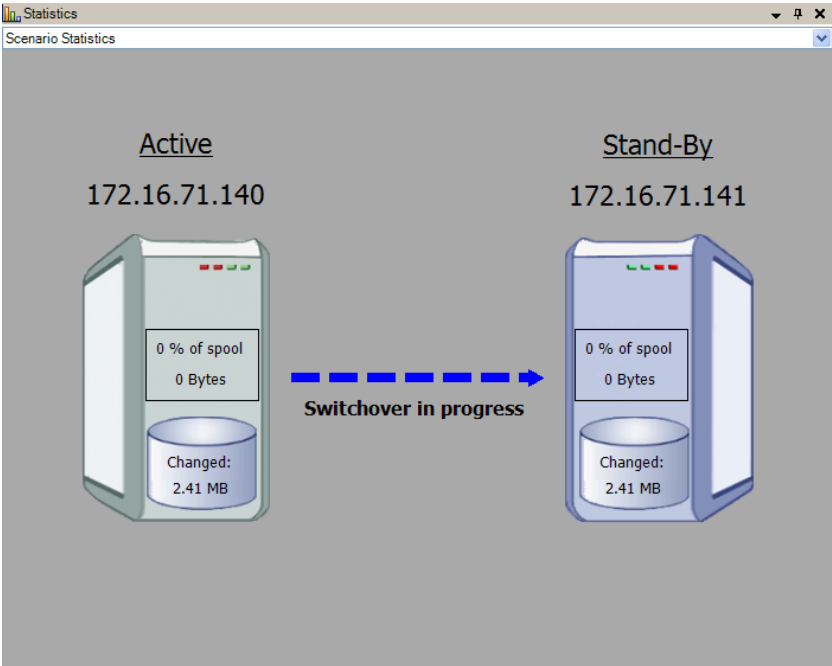
Note: Though the following steps show Exchange scenario screens as examples, the procedure is similar for all server types.

To initiate manual switchover

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

3. Click OK.

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



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

After the switchover is completed the scenario stops.

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

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

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

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

Initiate Switchback

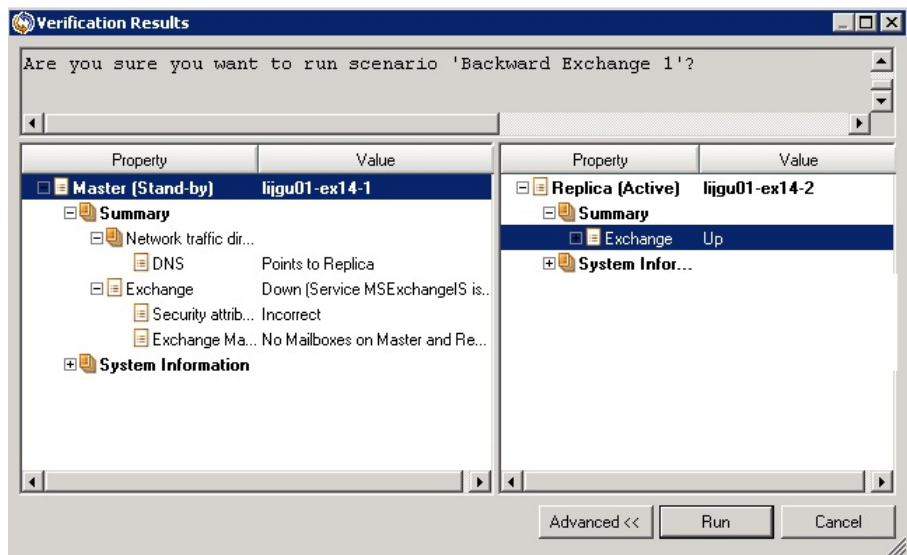
After a switchover is initiated, whether manually or automatically, at some point, you will want to reverse the server roles and make the original Master the active server and the Replica the standby server. Before you switch back the roles between servers, decide if you want the data on the original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

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

To initiate manual switchback

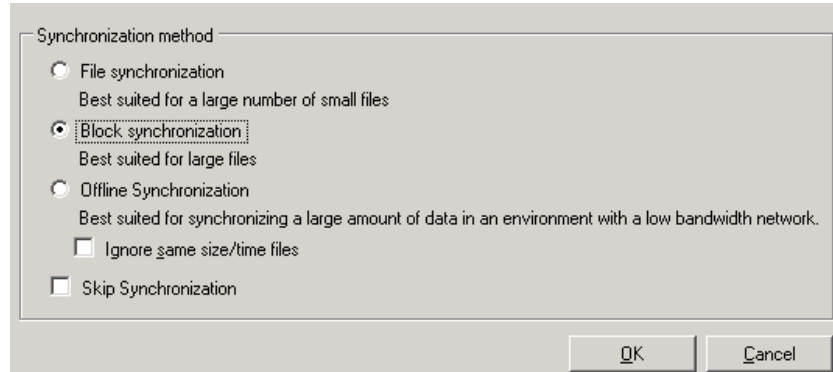
1. Ensure that both Master and Replica servers are available on the network and that the Engine is running.
2. Open the Manager and select the desired scenario from the Scenario pane.
3. Perform one of the following:
 - If the scenario is already running, skip directly to Step 4
 - If the scenario is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

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

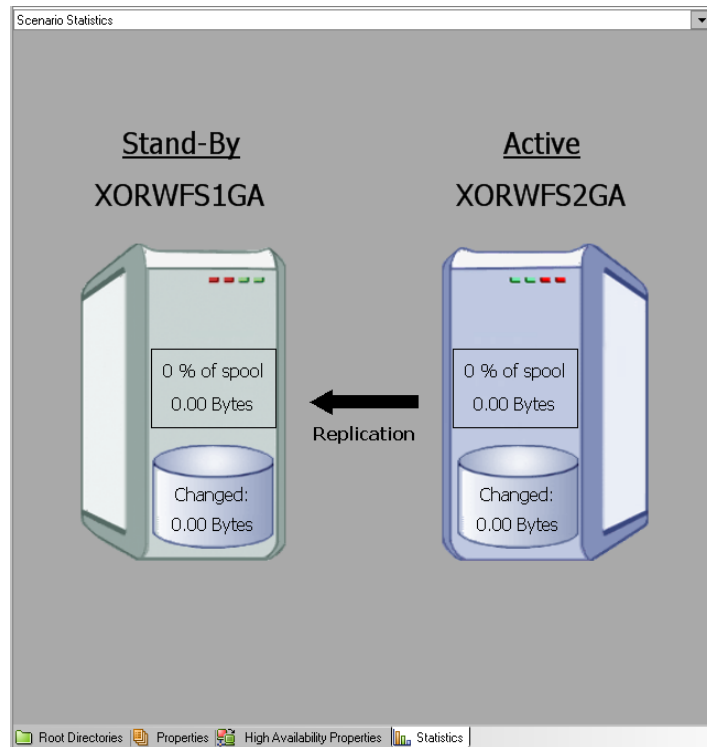


- b. Select a synchronization method from the Run dialog and click OK to start resynchronization.

Note: See the *Arcserve RHA Administration Guide* for more information on Synchronization Methods.



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



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

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

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

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

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

Switchover Considerations


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

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

Run a Scenario

You can run a single scenario using the following procedure:

To run the scenario

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

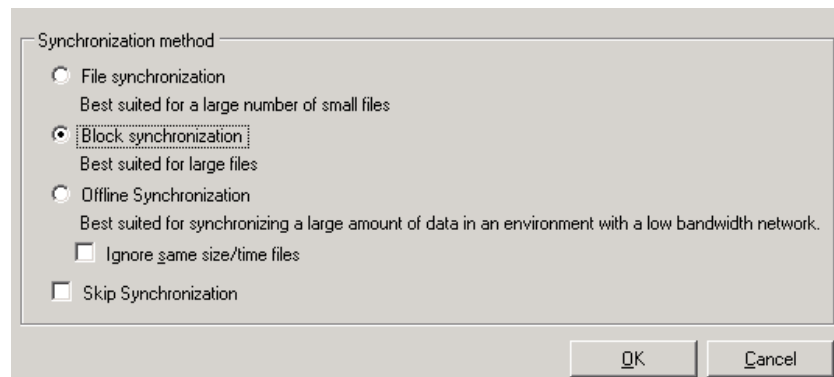
Before initiating synchronization and replication, Arcserve RHA verifies your scenario configuration. When verification completes successfully, the Manager displays the message: *Are you sure you want to run scenario "scenario_name?"* If problems are discovered, the top pane displays any warning and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

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

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

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

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




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

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

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

Stop a Scenario

To stop a scenario

1. From the Scenario pane, select the scenario you want to stop.
2. To stop the scenario, click the Stop  button on the Standard toolbar.
A confirmation message appears prompting you to approve the scenario stopping.
3. Click Yes in the confirmation message. The scenario stops.

After stopping the scenario, the Manager no longer shows the green play symbol to the left of the scenario, the scenario state turns to Stopped by user and the Statistics tab is no longer available on the Framework pane.

View a Report

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

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

To view reports

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

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

The Report Center opens in a new window.

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

Available Reports per Scenario

Scenario Name	Synchronizatio	Difference	Replication	Assessment Mode	Assured Recovery	Total Reports	
Exchange 1	1	0	0	0	0	1	

Reports

Drag a column header here to group by that column

Host	Changes	Date	Time	Type	Summary	Detailed	Size (bytes)
------	---------	------	------	------	---------	----------	--------------

The Report Center consists of two tables:

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

Host	Changes	Date	Time	Type	Summary	Detailed	Size (bytes)
XORVMEXCH267-1	Unknown	Today	03:29:37	Assured Recovery			811
XORVMEXCH267-1	Changes found	12/07/06	22:28:48	Synchronization			28415

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

The report you selected appears in a new window.

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

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

Chapter 3: Protecting Hyper-V Environments

This section contains the following topics:

[Hyper-V Server Configuration Requirements](#) (see page 59)

[Hyper-V Replication and High Availability](#) (see page 68)

[Switching Over and Switching Back](#) (see page 75)

[The Data Recovery Process](#) (see page 84)

[Additional Information and Tips](#) (see page 88)

[Manage Services](#) (see page 91)

Hyper-V Server Configuration Requirements

- Configure the same number and type of network connections on the Replica that exists on the Master.
- (In the Active Directory environment) Both Master and Replica servers should reside in the same Active Directory forest and should be members of the same domain or trusted domains.
- Make sure to change the default location of the virtual machine. You can change the location to any other custom location.
- For supported Hyper-V versions, see the Arcserve RHA Compatibility Matrix on Arcserve Support.

Hyper-V HA Configuration

Because Hyper-V is a Windows Server feature, you must set up two Windows Server 2008 machines, one Master and one Replica, to enable Arcserve RHA. You can use Arcserve RHA with only one Hyper-V server, but your protection will be limited to only Replication.

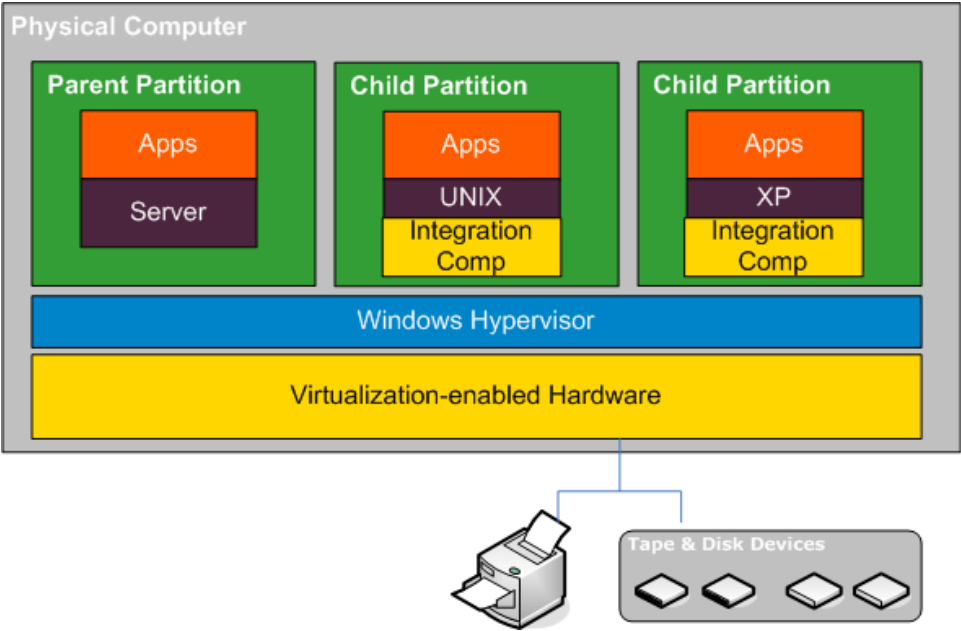
Note: For supported Hyper-V versions, see the Arcserve RHA Compatibility Matrix on Arcserve Support.

For HA, you must install the Hyper-V Integration Components on every Guest OS in your Hyper-V environment. This is so Arcserve HA can determine VM hostnames. For a list of supported guest operating systems, please see the Microsoft website.

Install Arcserve RHA on each Windows Server 2008 machine at the Server (Parent Partition) level. To achieve failover at the individual VM level, instead of merely at the Hyper-V Server, you must create scenarios for each VM in the environment. Arcserve RHA permits you to select more than one VM while creating the scenario, and then splits the entire selection into "sub-scenarios", effectively assigning one VM to a scenario. The names of each scenario are based on the names of their respective virtual machines.

In the following image, two virtual machines have been set up on a Hyper-V system. Integration Components were installed in each guest operating system. The Engine is installed on the Hyper-V machine, not in the VMs.

Note: Arcserve RHA supports virtual machines with Windows as a guest operating system.



Hyper-V VM Auto-Discovery

When you create a Hyper-V scenario, the Engine collects information about all virtual machines on the Hyper-V server and stores this data in a scenario configuration file. This is an .XML file. The following information is collected:

Virtual Machine Settings

These settings are stored in an XML file named <VM GUID>.XML, where VM GUID is the unique global identifier for the virtual machine. This file resides in the %ProgramData%\Microsoft\Windows\Hyper-V\Virtual Machines folder.

Virtual Machine Folder

There is a folder for each virtual machine and the name of the folder is <VM GUID>. This folder contains VSV and BIN files when a virtual machine is running.

Virtual Hard Disks (VHDs)

VHDs are hard disks associated with each VM. These files could be in different locations and have a .VHD extension.

Snapshots

All snapshots for a virtual machine are placed in the snapshot folder. All files inside a snapshot are considered for replication. Any VHD files in the snapshot have the extension, .AVHD. Multiple snapshots for a VM have multiple AVHD files. All files are discovered and replicated.

ACLs

In addition to files and folders, security data such as ACLs and owner information for each folder and file are replicated. This security information is needed to register and run the virtual machine on the Replica server.

Virtual Network Configuration

If a VM is connected to a virtual network, the network data setting is included in the XML configuration file.

Configure Hyper-V for High Availability

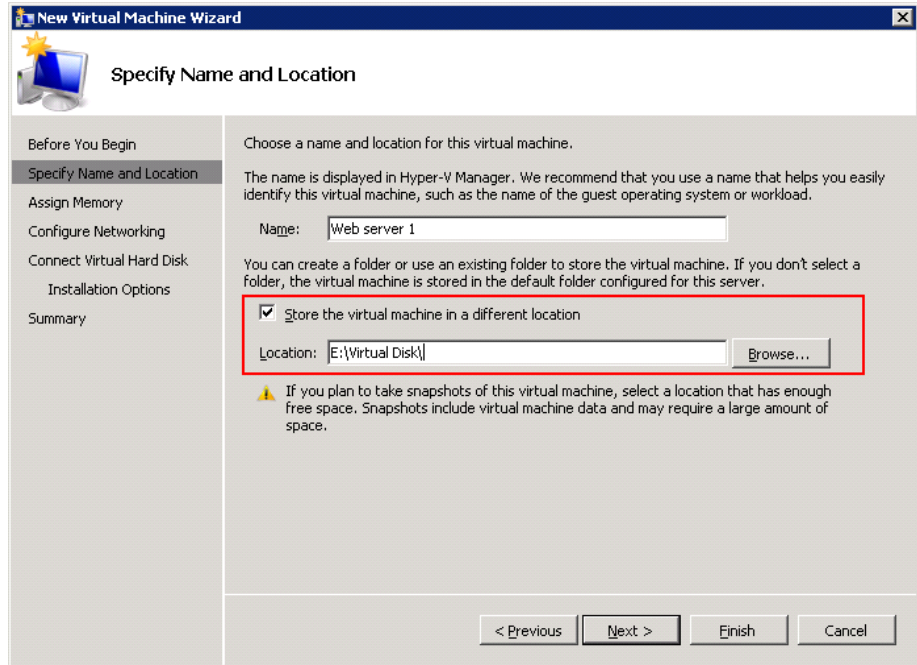
The following information applies only to Arcserve RHA High Availability. The default Hyper-V settings are usually sufficient when using Arcserve RHA but there are some default changes you should make to ensure virtual machines can easily and successfully be started on the Replica when the Master is unavailable:

- Integration Components are required by default. However, if you set the Is Alive, Check Method, Connect to DB property to Off, Arcserve RHA skips this check.
- The Snapshot folder cannot be overlapped for any Arcserve RHA-protected virtual machine but Microsoft puts all virtual machine snapshots in the same folder by default.

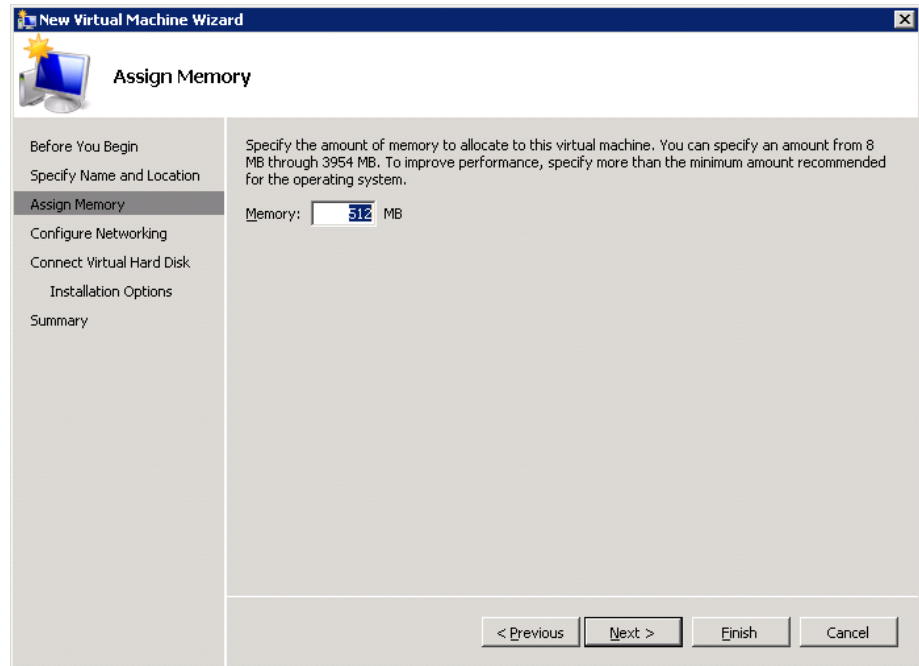
The following procedure uses a Windows guest as an example, but works for any supported Hyper-V guest operating system.

To configure Hyper-V for High Availability

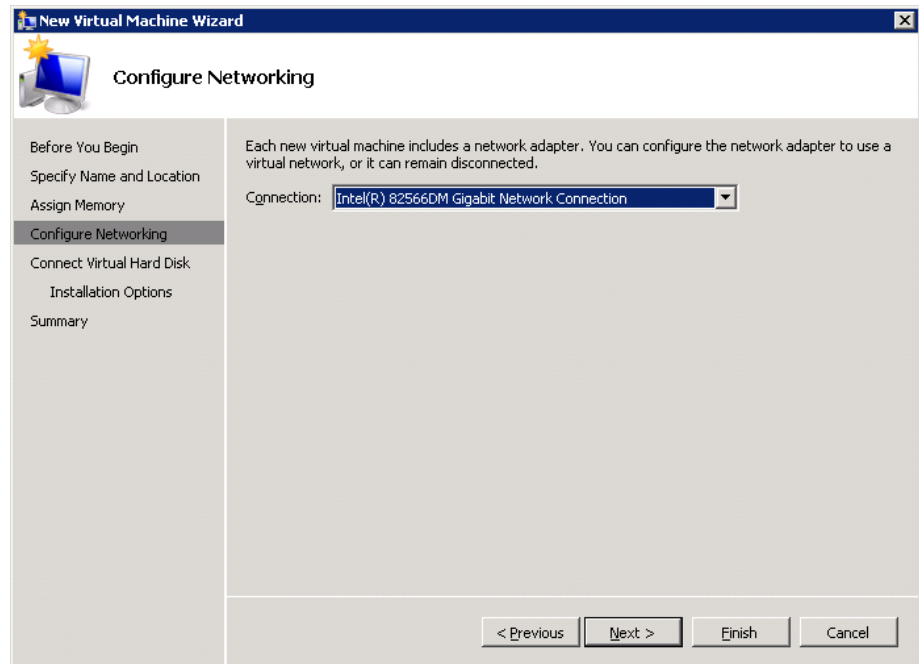
1. Launch the Hyper-V Manager wizard.
2. From the New Virtual Machine wizard dialog, enter a Name, select the **Store the virtual machine in a different location** option and then provide the location. We recommend changing the default path to a NAS or SAN volume. Click **Next**.



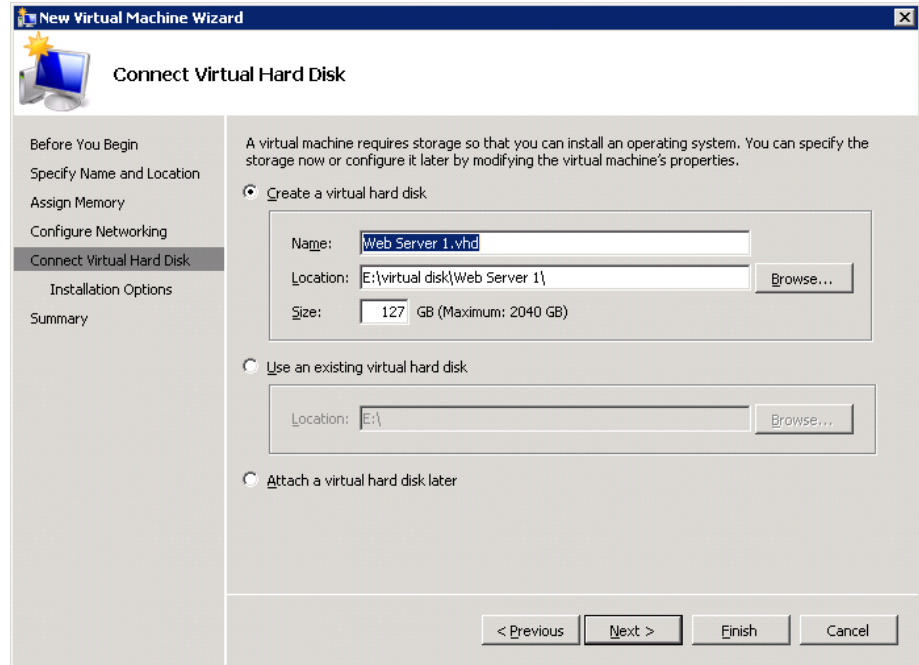
- From the Assign Memory dialog, set guest memory. Typically, 512 MB is sufficient for a single service guest. Click **Next**.



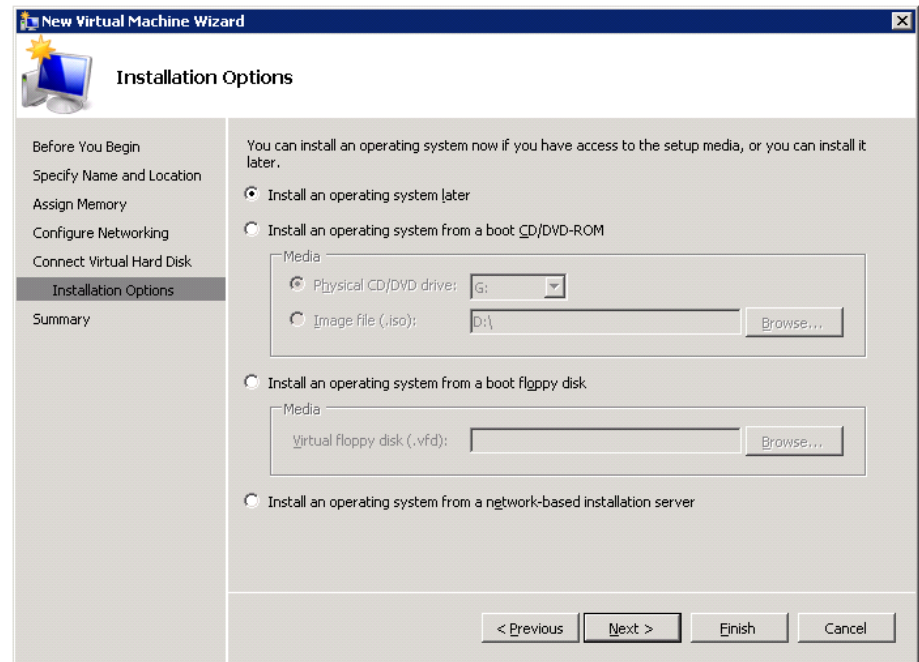
- From the Configure Networking dialog, select the virtual network adapter. Make sure you select a virtual network adapter that is bound to the physical network adapter. Click **Next**.



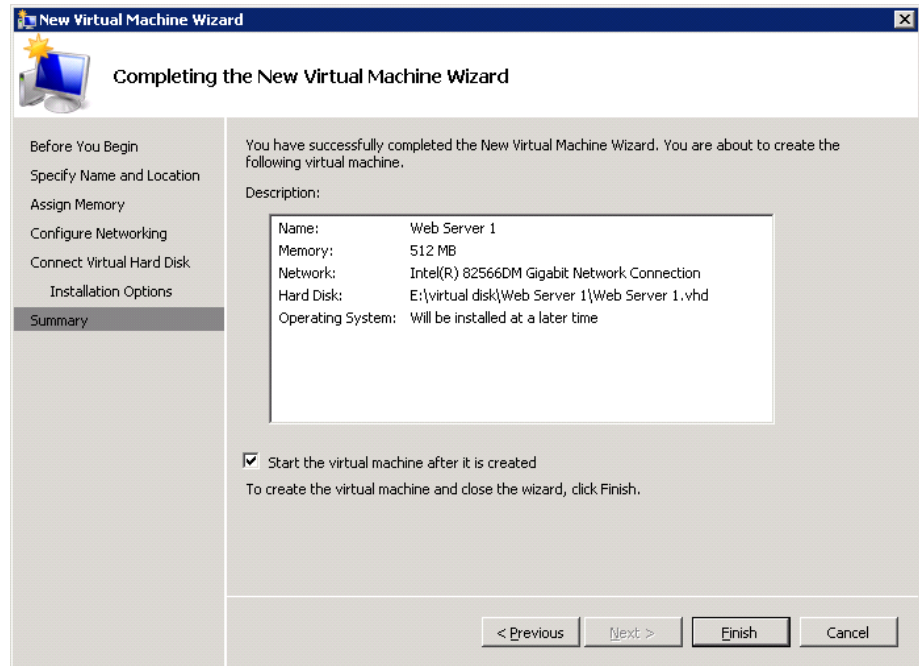
- From the Connect Virtual Hard Disk dialog, select **Create a virtual hard disk**. By default, the virtual disk is created under the VM folder. You may also use an existing disk. Click **Next**.



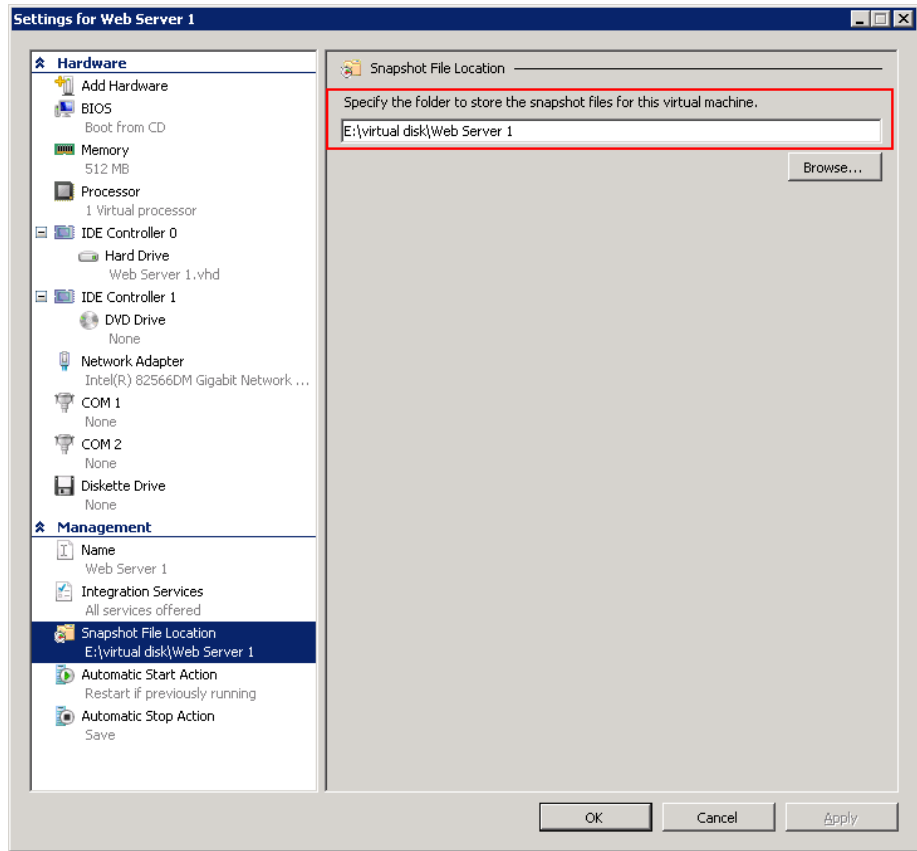
- From the Installation Options dialog, select **Install an operating system later** and click **Next**.



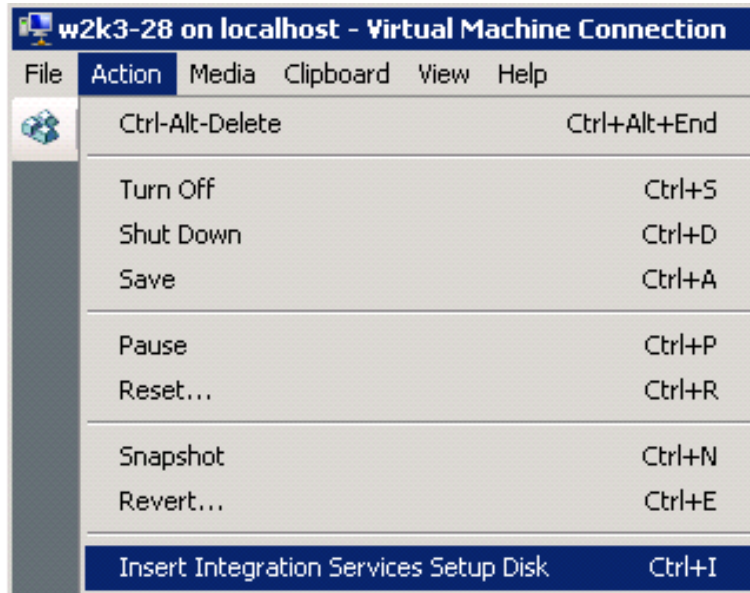
- From the Completing the New Virtual Machine Wizard dialog, select **Start the virtual machine after it is created** and click **Finish** to exit the wizard.



- 8. The snapshot folder is automatically set to the virtual machine folder. Click **OK**.



- 9. Once the operating system is ready, connect to the virtual machine. Click the Action menu and select **Insert Integration Services Setup Disk**.



10. Install the integration services in the guest OS.

You should also configure the Hyper-V Replica as follows:

- Install 64-bit Microsoft Windows Server 2008 with the Windows6.0KB950050-x86.msu Hyper-V patch on the Master server, with a CPU and motherboard capable of running Hyper-V. The Replica server can run any Windows-based system.
- Configure the same number of network connections on the Replica that exists on the Master.

Considerations for Hyper-V scenarios across WAN

This section describes how to create a Hyper-V scenario and ensure a smooth failover across different subnets in WAN environment. First set up an additional IP address and then specify the details while creating the Hyper-V scenario.

1. Set up an additional NIC on the VM running on the master Hyper-V server. The IP address of this NIC will be used on Replica after the VM fails over.

Note: Make sure that appropriate route table entries are added or updated to make both IP addresses and subnets work.

2. Create the Hyper-V scenario as usual. However, set the High Availability properties as described in the following steps:

- a. Expand Switchover and click Virtual Network mapping and choose the appropriate network mapping. Map the additional NIC/IP that you added. This IP address is used when the Replica server takes over.

- b. Enter the DNS server IP details.

Note: make sure DNS record can be updated by the RHA engine. See [Log On Account Conditions](#). (see page 8)

- c. Expand Network Traffic Redirection and click Redirect DNS. Select On and enter the following details:

Virtual Machine IPs on Master server in DNS

IP Address

Specifies the IP address of the virtual machine on the master server

Virtual Machine IPs on Replica server in DNS

IP Address

Specifies the additional IP address that you specified in step 1.

Hyper-V Replication and High Availability

Create a New Hyper-V Replication Scenario

Replication configuration properties are stored in scenarios. You need to create a scenario for each server you wish to protect.

Note: For Hyper-V HA/DR scenarios, RHA does not support the Hyper-V CSV (Cluster Shared Volume) clusters.

To create a Hyper V Replication Scenario

1. Open the Manager and choose Scenario, New or click the New Scenario button.
The Welcome dialog opens.
2. Choose **Create a New Scenario**, select a Group from the list and click **Next**.
Note: All scenarios you create are put in this scenario group. If you do not change the name, the final group name includes the Master server name as part of it.
3. The Select Server and Product Type dialog opens. Choose **Hyper-V, Replication and Data Recovery Scenario (DR)** and click **Next**.
4. The Master and Replica Hosts dialog opens. Select or type a Scenario group Name, enter the Hostname or IP Address and Port number for both Master and Replica servers, enable the **Verify Arcserve RHA Engine on Hosts** option and click **Next**.
5. Wait for Engine Verification to complete. Click **Install** if you need to upgrade the Engine service on one or both servers. When ready, click **Next**.
6. The Select Database for Replication dialog opens. Review the list of results auto-discovered on the Master server. By default, all VMs are replicated. Clear choices if desired and click **Next**.
7. The Replica Root Directories dialog opens. Accept the defaults or select the desired root directories on the Replica and click **Next**.
8. The Scenario Properties dialog opens. Set properties, as desired, and click **Next**.
9. The Master and Replica Properties dialog opens. Set properties, as desired, and click **Next**.
10. Wait for Scenario Verification to complete. Resolve any warning or errors and click **Next**.
11. From the Scenario Run dialog, click **Run Now** to initiate synchronization and activate the scenario or click **Finish** to run the scenario later.

Hyper-V Replication Properties

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

The Properties pane and its tabs are context-sensitive and change whenever you select a different node from a scenario folder. You must stop a scenario before configuring its properties. Certain values cannot be modified once set; they are noted. For full details on configuring scenario properties and their descriptions, see the *Arcserve RHA Administration Guide*.

Properties are organized into tabs on the Manager Framework pane. The tabs displayed are based upon server type, Arcserve RHA solution, and scenario status. Select the scenario for which you wish to change properties, and then select the appropriate tab. The following screen shows an example:

Scenario	State	Product	Server	Mode
w2k3copy1	Running	DR	HyperV	Online
w2k3copy2	Running	DR	HyperV	Online
w2k3copy3	Running	HA	HyperV	Online
2k832b01	Stopped ...	HA	HyperV	Online

ID	Sequence	Severity	Host/Scenario	Time	Event
IR00583	385	Info	172.16.78.112	1/13/2009 5:00:19 AM	Hyper-V synchronization bookmark created
SR00202	304	Significant	172.16.66.59	1/13/2009 4:07:48 AM	All modifications during synchronization period are replicated
IM00405	303	Info	w2k3copy3	1/13/2009 4:07:43 AM	Posting Synchronization report created at '01/13/09 04:07:42' to Reports
SR00120	301	Significant	172.16.66.59	1/13/2009 4:07:42 AM	Synchronization finished
IR00119	300	Info	172.16.66.59	1/13/2009 4:07:42 AM	Root directory d:\w2k3copy3\virtual machines synchronized
IR00119	299	Info	172.16.78.112	1/13/2009 4:07:33 AM	Root directory d:\w2k3copy3\snaps\hosts synchronized

Settings on the Root Directories tab

Select a Master Server from the Scenario Pane. Double-click its Directories folder to add or remove Master Root Directories. You cannot directly update the Hyper-V root directory. Double-clicking it launches Auto Discovery, which lists all virtual machines on the Master. You can add or remove virtual machines from the Auto Discovery Results dialog.

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

Settings on the Properties Tab

Scenario Properties

These settings establish default behavior for the entire scenario.

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

Master and Replica Properties

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

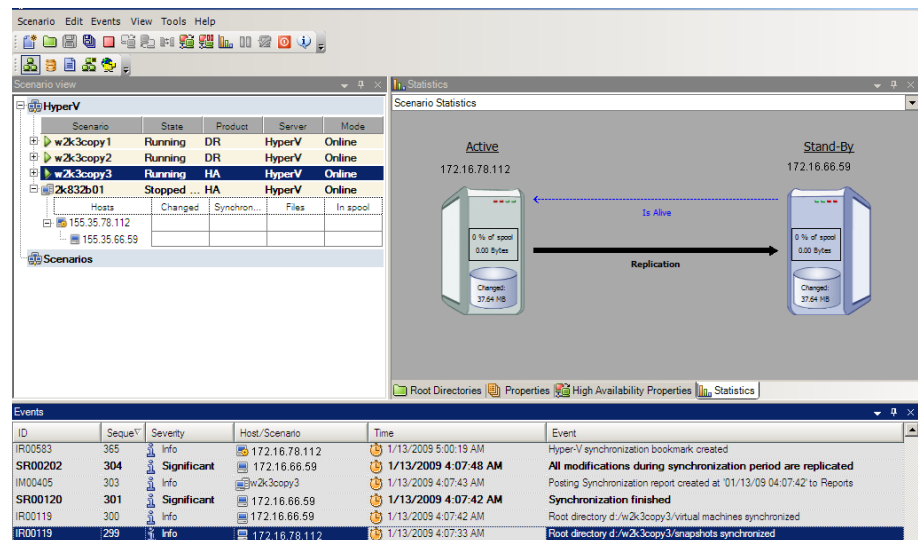
- Host connection properties -- Enter the IP address, Port number and Fully Qualified Name of the Master and Replica
- Replication properties -- Enable Hyper-V scheduled bookmarks for Master. These properties differ for Master and Replica. See the *Arcserve RHA Administration Guide* for more information.
- Spool properties -- Set the size, minimum disk free size and directory path. See [Spool Directory Settings](#) (see page 90) for more information.
- Event notification properties -- specify a script to run or choose email notification and write results to the event log.
- Report properties -- choose synchronization or replication reports, specify distribution or script execution.
- (Replica) Recovery properties -- set delay or data rewind properties.

Hyper-V HA Properties

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

The Properties pane and its tabs are context-sensitive and change whenever you select a different node from a scenario folder. You must stop a scenario before configuring its properties. Certain values cannot be modified once set; they are noted. For full details on configuring scenario properties and their descriptions, see the *Arcserve RHA Administration Guide*.

Properties are organized into tabs on the Manager Framework pane. The tabs displayed are based upon server type, Arcserve RHA solution, and scenario status. Select the scenario for which you wish to change properties, and then select the appropriate tab. The following screen shows an example:



Settings on the Root Directories tab

Select a Master Server from the Scenario Pane. Double-click its Directories folder to add or remove Master Root Directories. You cannot directly update the Hyper-V root directory. Double-clicking it launches Auto Discovery, which lists all virtual machines on the Master.

Select a Replica Server from the Scenario Pane. For each Master Root directory, you must specify a Replica Root directory. Double-click the Directories folder for the Replica server.

Settings on the Properties Tab

Scenario Properties

These settings establish default behavior for the entire scenario.

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

Master and Replica Properties

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

- Host connection properties -- Enter the IP address, Port number and Fully Qualified Name of the Master and Replica.
- Replication properties -- Enable Hyper-V scheduled bookmarks for master. These properties differ for Master and Replica. See the *Arcserve RHA Administration Guide* for more information.
- Spool properties -- Set the size, minimum disk free size and directory path. See [Spool Directory Settings](#) (see page 90) for more information.
- Event notification properties -- specify a script to run or choose email notification and write results to the event log.
- Report properties -- choose synchronization or replication reports, specify distribution or script execution.
- (Replica) Recovery properties -- set delay or data rewind properties. Data rewind is On by default.

Settings on the HA Properties Tab

These settings control how switchover and switchback are performed.

- Switchover properties -- choose automatic or manual switchover, provide switchover hostname, virtual network mappings, and reverse replication settings
- Hosts properties -- specify the Master and Replica Fully Qualified Name
- Is Alive properties -- set the heartbeat frequency and check method
- Action upon Success properties -- defines custom scripts and arguments for use

Create a New Hyper-V High Availability Scenario

For Hyper-V, it is possible to perform switchover at the individual VM level, rather than the whole Hyper-V server, which means you need a separate HA scenario for every virtual machine in the Hyper-V environment. To make scenario creation easier, you can select multiple VMs from the Master that inherit the same scenario properties. This 'multiple VM' scenario is then split into sub-scenarios so you can manage each virtual machine independently.

To create a new Hyper-V HA Scenario

1. Start Manager. Select File, Create, New Scenario or click the New Scenario button.

The Welcome dialog opens.

2. Click Create a New Scenario and type a Scenario Group Name, or select one from the list and then click Next.

Note: If you do not assign a Scenario Group Name, all scenarios you create default to the Scenarios Group. This group becomes part of the scenario name and updates automatically to Hyper-V after the scenario wizard is completed.

The Select Server and Product Type dialog opens.

3. Select MS Hyper-V, 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, enable the Verify Engine on Hosts option, and then click Next.

You may be prompted for user credentials. If so, enter the appropriate credentials and click OK.

If you enabled engine verification, the Engine Verification dialog opens.

5. Wait for verification to complete. Click Install to install the Engine on the specified hosts or click Next.

The Database for Replication dialog opens, listing the auto-discovered virtual machines found on the Master server you specified. By default, all virtual machines are selected for replication. For HA, the entire VM with all related files must be selected.

6. Select or clear virtual machines for replication and click Next.

The Scenario Properties dialog opens.

7. Change properties, as desired, and then click Next. For more information, see the *Arcserve RHA Administration Guide*.

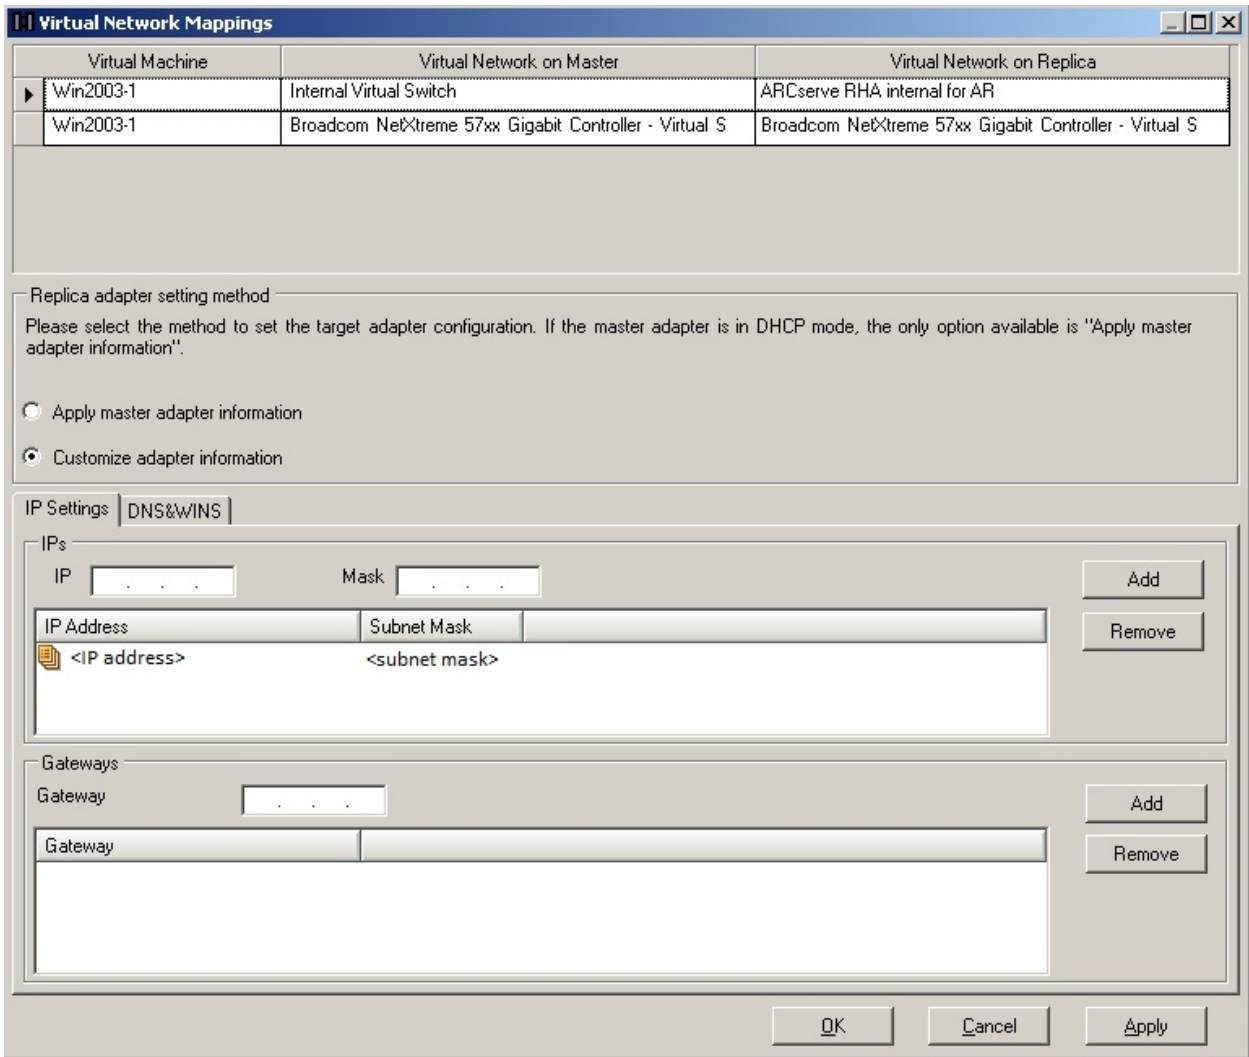
The Master and Replica Properties dialog opens.

- 8. Change properties, as desired, and then click Next. Scheduled Bookmarks are automatically set to On with a default of 1 hour, but you may adjust the schedule, if needed. For more information, see the *Arcserve RHA Administration Guide*.

Wait while the Switchover Properties dialog retrieves information.

- 9. When Switchover Properties opens, select Click to edit VN mappings.

The Virtual Network Mappings dialog opens.



10. Map the virtual machines listed to the desired virtual networks on the Replica and select the Replica adapter setting method.

Apply master adapter information

Specifies that the virtual machines on the replica use the same network adapter information as that of the master.

Customize adapter information

Specifies that you can manually enter the IP address, DNS, WINS, and Gateway information for the virtual machines.

Note: For virtual machines having boot volume on a dynamic disk, sometimes the network information does not configure successfully after switchover. You may have to manually set the IP address after the virtual machine starts up.

Limitation: If Hyper-v scenario protected VM sys volume is dynamic disk, some network mapping settings can't apply to the replica VM during Switchover process. RHA doesn't support network mapping for this type virtual machine.

11. Set any other switchover properties, as desired, and click Next.

The Switchover and Reverse Replication Initiation dialog opens. We recommend setting Switchover to Automatic and Reverse Replication to Manual.

12. Choose the desired method for initiating switchover and reverse replication and click Next.

Wait while the Scenario Verification process completes.

If Scenario Verification lists any errors, you must resolve them to continue. If any warnings are listed, you should also resolve them to successfully continue. After making changes, click Retry to repeat verification.

Click Next.

The Scenario Run dialog opens.

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

Switching Over and Switching Back

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

How Switchover and Switchback Work

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

- **Ping** -- a request sent to the VM on the Master to verify that the VM is up and responding. If the Is Alive check returns false, the Replica brings up the VM automatically.
- **User-defined check** -- a custom request you can tailor to monitor specific applications.
- **Connect to database** -- a request sent to the Master to verify that the appropriate services are running, and the VM is active and responding to pings.

These checks are performed sequentially (if enabled). 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 VM on the Master server is considered to be down. Then, depending on the HA scenario configuration, Arcserve RHA sends you an alert or automatically initiates a switchover.

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

- If you selected the Initiate Switchover manually option from the Switchover and Reverse Replication Initiation page, perform a manual switchover. For more information, refer to the topic, [Initiate Switchover](#) (see page 77).
- If you selected the Initiate Switchover automatically option, you can still perform a manual switchover, even if the Master is alive. You can initiate switchover when you want to test your system, or you want to use the Replica server to continue the application service while some form of maintenance is performed on the Master server. Triggered (automatic) switchover is in all ways identical to manual switchover performed by the administrator, except it is triggered by a resource failure on the master server rather than by an administrator manually initiating the switchover by clicking the Perform Switchover button. Server ping response, application service status, and VM connectivity are monitored. The timeout parameters are configurable and are more extensively covered in the *Arcserve RHA Administration Guide*.

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

- If you selected the Initiate Reverse Replication automatically option from the Switchover and Reverse Replication Initiation page, replication in the reverse direction (from Replica to Master) automatically begins after a switchover, one the original Master server becomes available again.
- If you selected the Initiate Reverse Replication manually option, you need to perform switchback manually. If you select the manual option and do not initiate a manual switchback, you must resynchronize data from Replica to Master, even after testing a clean switchover without a Master failure.


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

Initiate Switchover

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

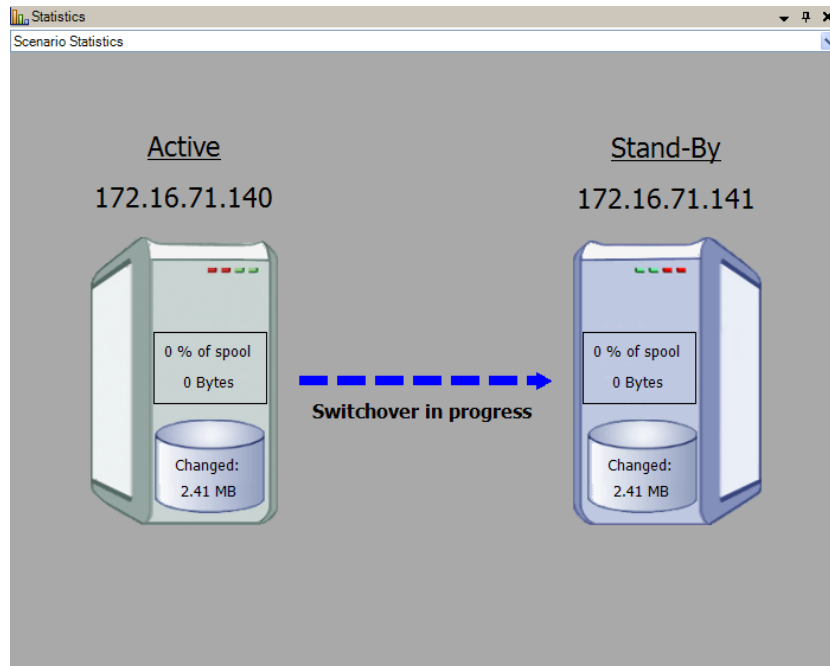
To initiate manual switchover

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



A confirmation message appears.

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



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

- After the switchover is completed the scenario stops:

HyperV XORWHYPERV1				
Scenario	State	Product	Server	Mode
XORWHypVM1	Editing	DR	HyperV	Online
XORWHypVM2	Editing	HA	HyperV	Online
XORWHypVM1 1	Stopped by user	DR	HyperV	Online
Hosts				
	Changed	Synchronized	Files	In spool
XORWHYPERV1				
XORWHYPERV2				

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

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

Now, the Master becomes passive and the Replica becomes active.

Initiate Switchback


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

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

To initiate manual switchback

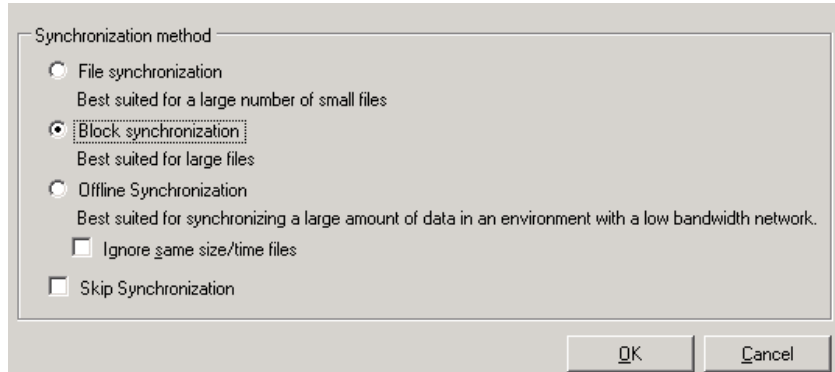
1. Ensure that both Master and Replica servers are available on the network and that the Engine is running.
2. Open the Manager and select the desired scenario from the Scenario pane.
3. Perform one of the following:
 - If the scenario is already running, skip directly to Step 4
 - If the scenario is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

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

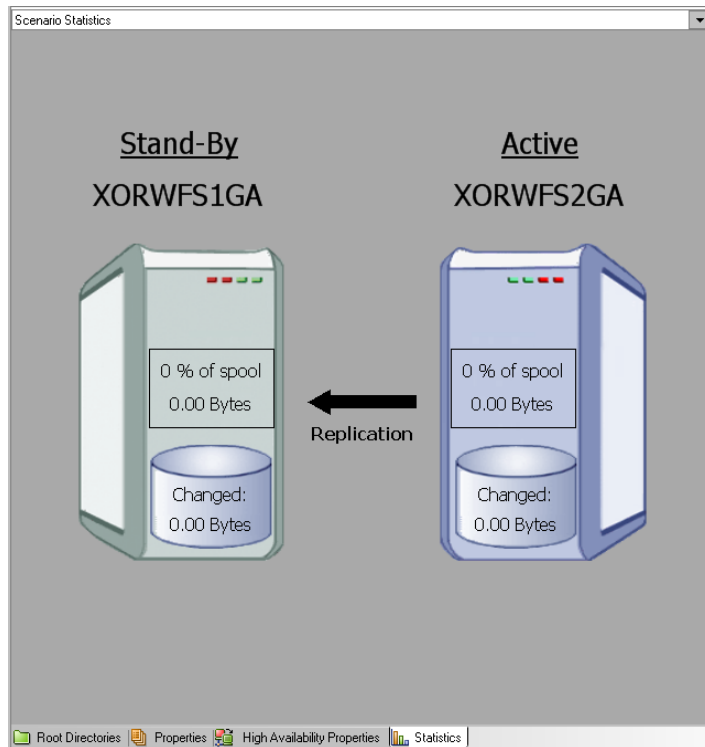


- b. Select a synchronization method from the Run dialog and click OK to start resynchronization.

Note: See the *Arcserve RHA Administration Guide* for more information on Synchronization Methods.



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



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

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

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

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

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

Switchover Considerations

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


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

Run a Scenario from Outside the Wizard

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

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

To start a scenario:

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

Arcserve RHA verifies the scenario before running it.

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


Notes:

- 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.

When no error is reported, the **Run** dialog appears.

3. From the Run screen, select the following and then click OK:
 - **Synchronization Method** -- For database and virtual machine applications, Block Synchronization is usually best but for File Servers or other applications with large numbers of small files, choose File Synchronization. See the *Arcserve RHA Administration Guide* for more information.
 - **Ignore same size/time files** -- Disable this option for database applications. Enable this option for File Server applications to speed up the comparison process and reduce overall synchronization time. See the *Arcserve RHA Administration Guide* for more information.
 - **Skip Synchronization** -- Select this option only if you are certain that the data in the Master and Replica root directories is identical.

The Manager indicates that the scenario is running with a green play symbol to the left of the scenario, and with the scenario's state, which turns into **Running**:



Scenario	State	Product	Server	Mode
XORWHypVM1	Editing	DR	HyperV	Online
XORWHypVM2	Editing	HA	HyperV	Online
XORWHypVM1 1	Running	DR	HyperV	Online
Hosts				
	Changed	Synchronized	Files	In spool
XORWHYPERV1				
XORWHYPERV2				


Once a scenario is running, a Statistics tab appears at the bottom of the Framework pane, displaying a graphical view of the replication.

By default, once synchronization occurs, a Synchronization Report is generated. To view the report, refer to the topic, [View a Report](#) (see page 56).

Note: You can also generate a Replication Report on a regular basis to monitor the replication process on each participating server. For more information, see *Arcserve RHA Administration Guide*.

Stop a Scenario

To stop a scenario

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

View a Report

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

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







To view a report

1. To view a report, first you need to open the Report Center. There are two ways to open it:
 - On the Overview Page, click the **Report Center** link on the **Quick Start** pane on the left.
 - 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:

Reports							
Drag a column header here to group by that column							
Host	Changes	Date	Time	Type	Summary	Detailed	Size (bytes)
XORWHYPERV2	Changes found	12/18/08	22:04:07	Synchronization			1923
XORWHYPERV2	Changes found	12/18/08	04:49:54	Synchronization			2548
XORWHYPERV2	Changes found	12/18/08	04:06:35	Synchronization			3479

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

The report you selected appears in a new window.

The Data Recovery Process

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

Setting Bookmarks

A *bookmark* is a checkpoint that is manually set to mark a state that you may want to rewind back to. We recommend setting a bookmark just before any activity that may 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**, in the Replica Properties list.
- You cannot set bookmarks during the synchronization process.
- Hyper-V HA allows you to schedule bookmarks.

To set a bookmark

1. When the required scenario is running, select the Replica host from which you want to rewind data.
2. From the menu, select **Tools, Set Rewind Bookmark**.

The **Rewind bookmark** dialog appears:



The text that appears in the **Rewind bookmark** dialog will appear 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. It is recommended to give a meaningful name that will later help you recognize the required bookmark. Then, click OK.

The bookmark is set.

In Master and Replica Properties, the Enable scheduled bookmark option is Off by default. Scheduling bookmarks allows you to use the Data Rewind option (default is Off). Set a regular bookmark schedule by clicking the Value column for Schedule Setting.

How to Restore Data on Hyper-V Machines

The process of restoring data from a Hyper-V virtual machine is much like restoring data in any other scenario with the following conditions:

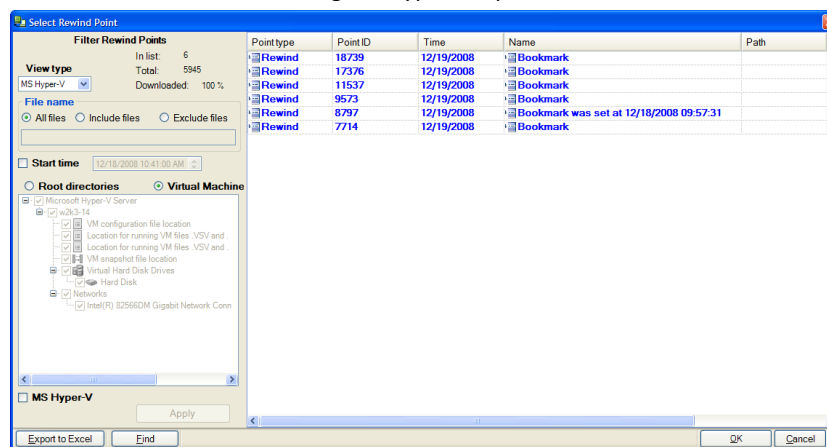
- **Stop the VM** -- Arcserve RHA automatically powers off the VM before recovery so that the existing VM can be overwritten. After restore completes, you need to restart the VM manually.
- **Select a bookmark** -- You must roll data back to a specific point in time called a bookmark, or rewind point. The default frequency is 1 hour, but bookmarks can be set at frequencies you define. When you create the Hyper-V scenario, ensure the Enable Hyper-V Scheduled Bookmarks setting is set to On from the Master and Replica Properties screen. For existing scenarios, you can edit this property manually.
- **Synchronize data** -- Use File or Block synchronization.

Recover Hyper-V Data with Rewind Points

For Hyper-V scenarios, the Data Rewind recovery method also applies.

To recover lost data using rewind points in a Hyper-V Server HA scenario

1. From the Manager, select the desired scenario and stop it.
2. From the Manager, select the Replica host to enable Restore Data options.
3. From the Tools menu, select Restore Data, or click the Restore Data button to open the Recovery Method dialog.
4. Choose the desired Rewind Data method, depending on whether you want the rewind data synchronized back to the Master or left on the Replica only. When you choose a Rewind Data option, a Recovery Scenario is automatically created. This Recovery Scenario runs until the end of the rewind process. Click Next to continue.
5. Wait while the Rewind Point Selection dialog retrieves information. When the Select Rewind Point button is enabled, click it to continue.
6. The Select Rewind Point dialog for Hyper-V opens.



This dialog displays information specific to Hyper-V scenarios. You can choose to display rewind points by file name or start time. Select the desired rewind point and click OK to return to the Rewind Point Selection dialog, which now displays the rewind point you selected.

7. Click Next to open the Synchronization Method dialog.
8. Choose Block Synchronization and then click Finish.

Arcserve RHA rewinds the data to the point you selected. After the rewind process ends, the following message is displayed 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 completed, the temporary Recovery Scenario is stopped and then deleted. If you wish, you can view the Synchronization Report that is generated by default. At this time, Replication can restart on the original scenario.

How to Start a Hyper-V VM on the Replica Server

When a virtual machine fails over to the Replica server, Arcserve RHA starts the virtual machine on the Replica automatically.

Additional Information and Tips

This section provides you with helpful information concerning the application.

- By default, the spool is located in the Arcserve RHA installation /tmp directory. You can change the default location by modifying the pathname for spool directory. It is best to configure the spool on a non-SQL database or log file drive. Using a dedicated volume for the spool folder can increase performance under high load. If you do change the spool location, please remember to remove the new path from the anti-virus scans, both scheduled and real-time.
- Arcserve RHA supports bandwidth limitation and bandwidth limitation scheduling. If you require such features, please consult the *Arcserve RHA Administration Guide*.

Troubleshooting Hyper-V

The following information is provided to help you resolve certain problems:

CV01378 Hyper-V HA scenario has unassigned virtual network mapping

Reason:

The scenario you created discovered more than one virtual network on the Replica. You must map additional virtual networks manually.

Action:

From High Availability Properties, expand Switchover properties and click "Click to edit virtual network mapping" to map virtual networks manually.

Could not edit virtual network mappings

Reason:

This is an internal error.

Action:

1. Restart the Control Service.
2. Re-create the scenario.
3. Collect log messages and the scenario .xmc file.
4. Contact Support.

Could not retrieve list of virtual networks from replica

Reason:

This is an internal communication error, but the scenario is likely correctly configured.

Action:

1. Collect log messages and the scenario .xmc file.
2. Contact Support.

The Replica Server has no virtual networks defined. This could mean the Replica has no Hyper-V role enabled or no virtual networks were configured.

Reason:

The Replica Server has no virtual networks defined. This could mean the Replica has no Hyper-V role enabled or no virtual networks were configured.

Action:

Ensure the Hyper-V role is enabled on the Replica server. Ensure at least one virtual network is defined.

Editing error

Reason:

You left a required field blank or provided an invalid entry.

Action:

Ensure all fields are completed and entries are valid.

Spool Directory Settings

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

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

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

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

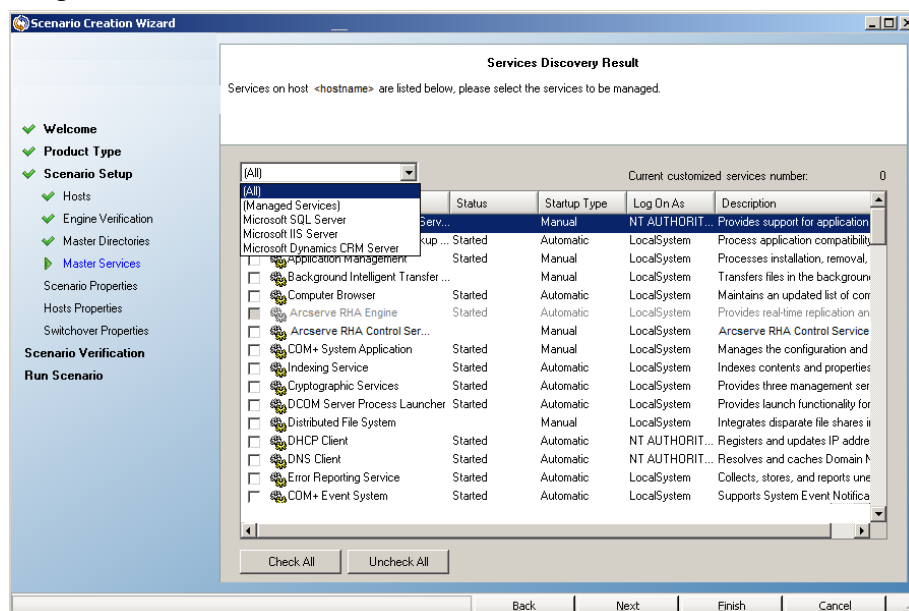
Manage Services

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

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

The following steps are for Custom Application scenarios.

To manage services



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


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

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

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

Services Setting

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



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

Chapter 4: Protecting Full Systems

Full System scenarios let you protect or move a physical system to a virtual machine. With Full System scenarios, you can switch a physical machine over to a virtual server manually or automatically, perform assured recovery testing, or set bookmarks that can be used as restore points to either the original machine, or to a new machine.

Arcserve RHA supports dynamic disks including Simple, Striped, Mirrored, Spanned, and RAID-5 volumes. You can have both master and replica as dynamic disks. Dynamic disks let you create volumes that span multiple disks.

Note: When you create a bookmark, RHA engine calls the application VSS writer to create an application consistent snapshot on the master. RHA has four parameters (EnableVSSWriters, ExcludedWriterList, IncludedWriterList, StartWriterServicesList) in the ws_rep.cfg file to customize how the RHA engine calls the application VSS writer to create the snapshot.

This section contains the following topics:

[How Full System Scenarios Work](#) (see page 93)

[Unified Extensible Firmware Interface \(UEFI\) Support for Full System Scenarios](#) (see page 96)

How Full System Scenarios Work

Use Full System scenarios to protect any server, regardless of application type. Replication, High Availability and Assured Recovery are supported. Full System scenarios are application-independent and let you transfer an entire physical machine (Master), including all system status information, to a virtualized Replica that supports the guest operating system of the active server.

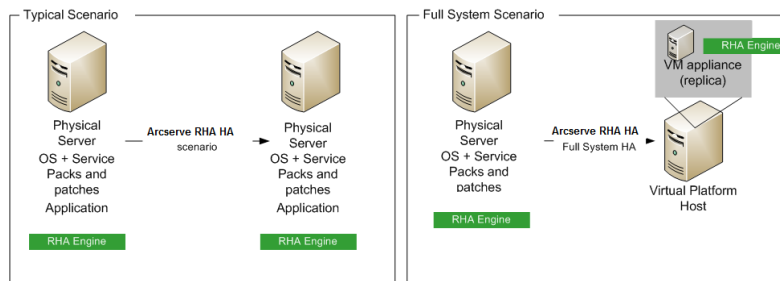
The following virtualized environments are supported as the Replica server in Full System scenarios:

- Microsoft Hyper-V
- Citrix XenServer (Xen)
Note: Review XenServer Configuration Limitations before using XenServer on www.citrix.com.
- ESXi/vCenter Server
- Amazon EC2

Note: Refer to the Arcserve RHA Release Notes or the Compatibility Matrix on Arcserve Support for supported versions.

Data from a physical Master is replicated *not* to a physical Replica, as it is in typical Arcserve RHA scenarios, but to a virtual hard drive file stored on a VM. This virtual machine, called an *appliance* in Full System scenarios, runs on a virtual platform host and is configured with the Arcserve RHA Engine. At switchover, the physical Master is disabled, a new virtual machine is created on the virtual platform host and the virtual disk file is mounted as a disk. The new VM is booted and integration services are deployed.

Note: You cannot run two Full System scenarios that protect the same physical Master.



The following list highlights how primary Arcserve RHA operations work with Full System scenarios.

Scenario Creation -- Create a Full System scenario that identifies a physical machine as the Master server and a virtual machine as the Appliance (Replica) and the virtual platform that hosts this Appliance VM. Specify information unique to Full System scenarios, depending upon the virtual platform type:

- Volume List -- You can select which volumes on the Master to include in the scenario. The system and boot volumes must be protected and are always selected by default.
- Root Directory List -- Since there is no root directory concept in Full System scenarios, this list is used to guide which data is replicated to virtual disk.
- Local Directory on VM -- This is the location of the virtual disk stored on the virtual machine server. All virtual disks are stored here. It also stores the virtual disk volume mount point. Arcserve RHA replicates data from the source directory on the physical Master to the mount point on the virtual Replica.
- Network Adapter List -- This stores all network adapter information (adapter name, adapter ID, IP address, MAC address, and so on) on the physical machine. This information is used to create network mapping lists between the physical machine's adapter and the virtual machine's adapter.

- Network Mapping List -- When the virtual machine is created, the network adapter on the VM is set according to the information in this list.
- CPU/Memory Value -- When the virtual machine is created, Arcserve RHA sets the CPU number and memory size of the VM according to configuration properties defined in the scenario.
- Resource Pool -- For vCenter or ESXi platforms, choose a resource pool that specifies the shared resources for the VM.
- Datastore -- For vCenter or ESXi platforms, choose a datastore that specifies where the VM files are to be saved.

Scenario Start -- Arcserve RHA validates the stored scenario to verify no errors exist that could prevent a successful switchover, if needed.

Synchronization -- The physical Master creates and transfers the virtual disk file to the local directory on the virtual Replica. After transfer, the virtual disk file is stored in the directory you specified on the Root Replica Directory dialog during scenario creation.

Replication -- The physical Master replicates any data changes to the virtual disk file stored on the Replica in real time mode. The virtual disk file is mounted as a directory on the Hyper-V server file system.

Switchover -- If the physical Master is unavailable, Arcserve RHA initiates switchover according to the properties set during scenario creation. Arcserve RHA stops the Full System scenario and disables the physical machine network. The virtual machine is created on the server using the same name as the physical Master. The virtual disk is mounted and the network adapter for the virtual machine is added according to the network mapping list defined in the stored scenario. Finally, the virtual machine is booted. For end users accessing the failed physical Master, switchover takes several minutes.

Data Recovery -- Launch the Data Recovery Wizard and select a restore point. You can select the volumes to be restored and the target to which these volumes should be restored. Arcserve RHA creates and starts the restore scenario.

Assured Recovery -- You can perform manual or automatic assured recovery. Select the virtual Replica and click Replica Integration Testing, select the type of Assured Recovery and continue. Arcserve RHA stops applying journal changes, starts the AR process by creating a virtual machine with the specified virtual disks and then starts the VM. Arcserve RHA resumes applying journal changes when AR is stopped.

Unified Extensible Firmware Interface (UEFI) Support for Full System Scenarios

The Unified Extensible Firmware Interface (UEFI) is a specification definition for an interface between an operating system and firmware of the virtual platform. For more information about UEFI, visit <http://www.uefi.org/home/>.

Arcserve RHA supports selecting a computer with UEFI firmware as the master server. However, review the following considerations:

- When you use a machine with the UEFI firmware as the master server to create a Full System HA or AR scenario, then the virtual platform must also support booting from UEFI firmware, such as ESXi 5.1.
- When you use a machine with the UEFI firmware as the master server to create a Full System DR scenario and the virtual platform does not support booting from the UEFI file system (for example, ESX 4.0), then the start VM operation fails.

Index

A

- About This Guide • 7
- Add Arcserve-IP to Existing Scenarios • 44
- Add Arcserve-IP to New Scenarios • 46
- Add IP on the Master Server • 42
- Additional Information and Tips • 88
- Arcserve Product References • 3
- Automatic Redirection Using Switch Computer Name • 47

C

- Configure Hyper-V for High Availability • 61
- Configure the vCenter Server Database Server on the Replica • 15, 16
- Configure the vCenter Server License Server on the Replica • 15, 18
- Configure the vCenter Server Master Server • 14
- Configure the vCenter Server Replica Server • 14
- Considerations for Hyper-V scenarios across WAN • 67
- Contact Arcserve • 3
- Create a New Hyper-V High Availability Scenario • 73
- Create a New Hyper-V Replication Scenario • 68
- Create vCenter Server HA Scenarios • 21
- Create vCenter Server Replication Scenarios • 19

D

- Data Rewind • 29
- DNS Redirection • 42
- Documentation Changes • 4

H

- How Full System Scenarios Work • 93
- How Redirection Works • 41
- How Switchover and Switchback Work • 49, 76
- How to Restore Data on Hyper-V Machines • 86
- How to Restore Data on vCenter Machines • 31
- How to Start a Hyper-V VM on the Replica Server • 88
- Hyper-V HA Configuration • 60
- Hyper-V HA Properties • 70
- Hyper-V Replication and High Availability • 68
- Hyper-V Replication Properties • 69

- Hyper-V Server Configuration Requirements • 59
- Hyper-V VM Auto-Discovery • 61

I

- Initiate Switchback • 52, 79
- Initiate Switchover • 49, 50, 76, 77
- Introduction • 7

L

- Log On Account Conditions • 8, 67

M

- Manage Services • 91
- Manually Recover a Failed Server-Move IP Address • 31, 32
- Manually Recover a Failed Server-Switch Computer Name • 32
- Manually Recover Failed Server-IP and Switch Computer Name • 31, 33
- Move IP Redirection • 42

P

- Protecting Full Systems • 93
- Protecting Hyper-V Environments • 59
- Protecting VMware vCenter Server Environments • 11

R

- Recover Active Server • 27
- Recover Hyper-V Data with Rewind Points • 87
- Recover Lost Data from Replica • 26
- Recovering Servers • 27, 31
- Redirection Methods • 41
- Register Arcserve RHA Licenses • 8
- Related Documentation • 8
- Run a Scenario • 54
- Run a Scenario from Outside the Wizard • 82

S

- Scripts Redirection • 47
- Set Scenario Properties • 21, 22
- Setting Bookmarks • 28, 85
- Spool Directory Settings • 24, 70, 72, 90

Stop a Scenario • 56, 83
Supported Databases VMware vCenter Server • 17
Switch Computer Name Redirection • 46
Switching Over and Switching Back • 22, 48, 75
Switchover Considerations • 54, 81

T

The Data Recovery Process • 25, 84
Troubleshooting Hyper-V • 88
Troubleshooting Oracle Databases • 41
Troubleshooting vCenter Server Scenarios • 33

U

Unified Extensible Firmware Interface (UEFI) Support
for Full System Scenarios • 96

V

vCenter Server Replication and High Availability • 19
View a Report • 56, 83
Virtualization Considerations • 9
VMware vCenter Server RHA Configuration • 12