

CA ARCserve® Replication and High Availability

Virtualized Server Environments Operation Guide for Windows

r16



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 CA 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 CA. This Documentation is confidential and proprietary information of CA 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 CA governing your use of the CA software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and CA.

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 CA 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 CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA 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 CA 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 CA 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 CA.

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.

Copyright © 2012 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Product References

This document references the following CA Technologies products:

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

Throughout this Guide, the term, CA ARCserve RHA, is used to represent the entire product family, which was previously sold as CA XOssoft Replication (WANSync) and CA XOssoft High Availability (WANSyncHA).

Contact Technical Support

For your convenience, CA provides one site where you can access the information you need for your Home Office, Small Business, and Enterprise CA products. At arcserve.com/support, you can access the following:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Provide Feedback

If you have comments or questions about CA product documentation, you can send a message to techpubs@ca.com.

If you would like to provide feedback about CA product documentation, complete our short [customer survey](#), which is also available on the CA Support website, found at <http://ca.com/docs>.

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 9

About This Guide	9
Related Documentation	10
Log On Account Conditions	10
Register CA ARCserve RHA Licenses	10
Virtualization Considerations	11

Chapter 2: Protecting VMware vCenter Server Environments 13

VMware vCenter Server RHA Configuration	14
Configure the vCenter Server Master Server	16
Configure the vCenter Server Replica Server	16
vCenter Server Replication and High Availability	21
Create vCenter Server Replication Scenarios	21
Create vCenter Server HA Scenarios	23
Set Scenario Properties	24
The Data Recovery Process	27
Redirection Methods	44
How Redirection Works	44
DNS Redirection	45
Move IP Redirection	45
Switch Computer Name Redirection	48
Scripts Redirection	49
Switching Over and Switching Back	50
How Switchover and Switchback Work	51
Initiate Switchover	52
Initiate Switchback	54
Switchover Considerations	56
Run a Scenario	56
Stop a Scenario	58
View a Report	58

Chapter 3: Protecting Hyper-V Environments 61

Hyper-V Server Configuration Requirements	61
Hyper-V HA Configuration	62
Hyper-V VM Auto-Discovery	63
Configure Hyper-V for High Availability	64

Considerations for Hyper-V scenarios across WAN	69
Hyper-V Replication and High Availability.....	70
Create a New Hyper-V Replication Scenario	70
Hyper-V Replication Properties.....	71
Hyper-V HA Properties	72
Create a New Hyper-V High Availability Scenario	75
Switching Over and Switching Back.....	76
How Switchover and Switchback Work.....	77
Initiate Switchover	78
Initiate Switchback.....	80
Switchover Considerations.....	82
Run a Scenario from Outside the Wizard	83
Stop a Scenario.....	84
View a Report.....	84
The Data Recovery Process	86
Setting Bookmarks	86
How to Restore Data on Hyper-V Machines	87
Additional Information and Tips	89
Spool Directory Settings.....	91
Manage Services	92

Chapter 4: Protecting Full Systems 95

How Full System Scenarios Work	95
Configuring Master and Replica for Full System Scenarios	98
Configure HTTP Communication Protocol on vCenter Server 4.0 Systems	99
Create Full System Scenarios	99
Create Full System Scenarios for Hyper-V Platform.....	105
Configure the Web Proxy to Connect to the Cloud Service	108
Add a New Cloud Account.....	109
Change the Default AWS Region.....	111
Update Cloud Account Information	111
Delete a Cloud Account.....	111
Create a New EC2 Replica Instance.....	112
Start an EC2 Replica Instance	118
Stop an EC2 Replica Instance	118
Delete an EC2 Replica Instance	118
Create a New EC2 Data Replication Scenario.....	119
Create a New Full System EC2 High Availability Scenario	123
Additional Properties for Full System Scenarios	131
Configure Additional Properties in Full System-to-ESX Scenarios.....	132
Configure Additional Properties in Full System EC2 Scenarios	133

Redirection Methods for Full System Scenarios.....	135
Running a Full System HA Scenario	135
Operations on a Virtual Machine	136
Start or Stop a Virtual Machine.....	136
Delete Virtual Machine Resources	137
Restore Full Systems	138
Review Prerequisites.....	139
Create and Run the Recovery Scenario	140
DNS Redirection using the Update DNS Tool	143
Restoring Data to a Bare Metal Machine	144
Review Prerequisites.....	146
Prepare a Bare Metal Machine	146
Create and Run the Recovery Scenario.....	147
Verify Bare Metal Machine Properties.....	147
Restoring Data to a Bare Metal Machine after Failover.....	148
Review Prerequisites.....	150
Prepare a Bare Metal Machine	150
Create and Run the Recovery Scenario.....	151
(Optional) Perform a Manual Switchover	152
Verify Bare Metal Machine Properties.....	152
Perform Full System Assured Recovery Testing	153
Recover Active Server for Full System Scenarios	153
Additional Tools	154
Virtualization Hypervisor Login Setting Tool	155

Index

157

Chapter 1: Introduction

CA ARCserve Replication and High Availability (CA 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.

CA 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 CA 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 9)

[Related Documentation](#) (see page 10)

[Log On Account Conditions](#) (see page 10)

[Register CA ARCserve RHA Licenses](#) (see page 10)

[Virtualization Considerations](#) (see page 11)

About This Guide

This document describes how to implement a CA 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 *CA ARCserve RHA Installation Guide* and the *CA ARCserve RHA Administration Guide*.

Log On Account Conditions

The CA 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 CA ARCserve RHA Licenses

The CA 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 CA 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 CA 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 CA 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 CA 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, CA 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 14)

[vCenter Server Replication and High Availability](#) (see page 21)

[Redirection Methods](#) (see page 44)

[How Switchover and Switchback Work](#) (see page 51)

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 CA 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.

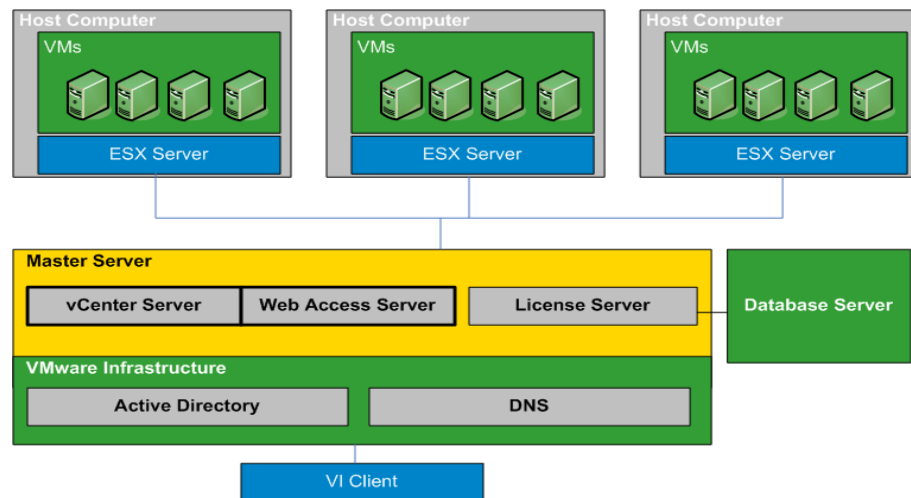
Important! Regardless of the method you choose, the vCenter Server and the Web Access Server must be installed on the same machine.

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 CA 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 CA 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 CA 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 CA 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 18)
- 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 20)

Note: To protect the distributed (remote) database servers, create separate CA ARCserve RHA scenarios for SQL or Oracle, as appropriate. However, if the License Server is deployed separately, CA 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

CA 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 CA 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 CA 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. CA 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 24) 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 50)

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 *CA ARCserve RHA Administration Guide*.

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

The screenshot shows the CA ARCserve RHA Manager Framework interface. The main window is divided into two panes: Scenarios and Properties.

Scenarios Pane: This pane displays a table of scenarios. The table has columns for Scenario, State, Product, Server, and Mode. The scenarios listed are Exchange - DR, Exchange - HA, Exchange - MovelP, and SQL. The Exchange - MovelP scenario is selected, and its details are shown in a sub-table below it.

Scenario	State	Product	Server	Mode
Exchange - DR	Editing	DR/AR	Exchange	Online
Exchange - HA	Editing	HA/AR	Exchange	Online
Exchange - MovelP	Stopped by...	HA/AR	Exchange	Online
SQL				

The sub-table for Exchange - MovelP shows a list of hosts and their status:

Hosts	Changed	Sent D...	Sent Fl...	Receiv...	Receiv...	In spool
192.168.1.13						
192.168.1.12						

Properties Pane: This pane shows the properties for the selected scenario. It has a table with columns for Property and Value. The properties are organized into tabs: General, Replication, Event Notification, and Report Handling.

Property	Value
General	
Replication	
Event Notification	
Report Handling	

Events Pane: This pane shows a list of events. The table has columns for ID, Sequence, Severity, Host/Scenario, and Time.

ID	Sequence	Severity	Host/Scenario	Time
SR00096	1785	Significant	192.168.1.13	5/2/2012 10:51:49 PM
SR00202	1784	Significant	192.168.1.12	5/2/2012 10:51:21 PM
IM00405	1783	Info	Exchange-MovelP	5/2/2012 10:51:20 PM
SR00120	1781	Significant	192.168.1.12	5/2/2012 10:51:19 PM
IR00119	1780	Info	192.168.1.12	5/2/2012 10:51:19 PM
IR00119	1779	Info	192.168.1.12	5/2/2012 10:50:55 PM

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 *CA 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 91) 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 *CA 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 CA 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.

CA 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 *CA ARCserve RHA Administration Guide*. Click **Finish**.

Once you finished initiating the recovery process, CA 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

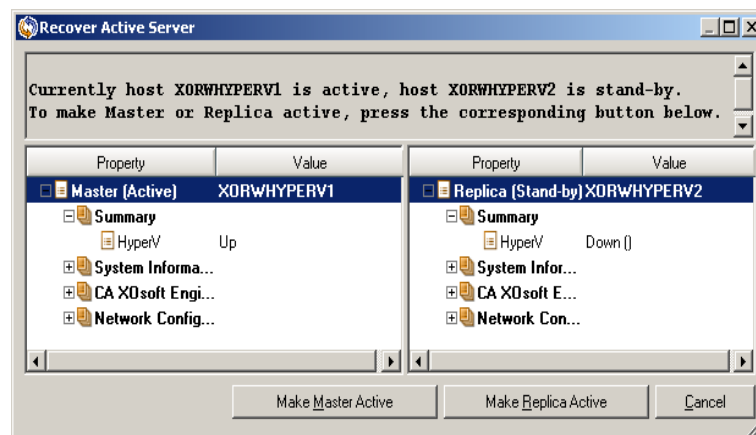
In certain circumstances, it may be necessary to forcibly make the Master or Replica server the active server without completing the data synchronization process.

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. CA 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, CA 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, CA 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 34).

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.

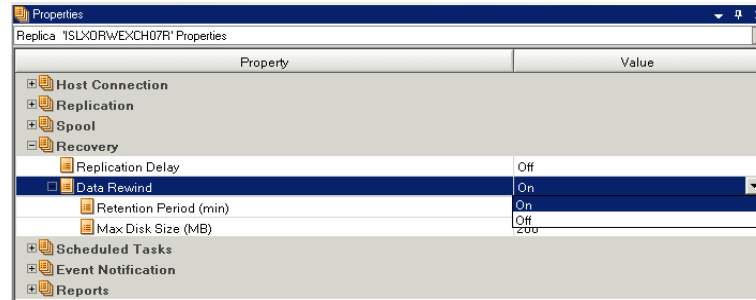
The bookmark is set.

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 *CA ARCserve RHA Administration Guide*.

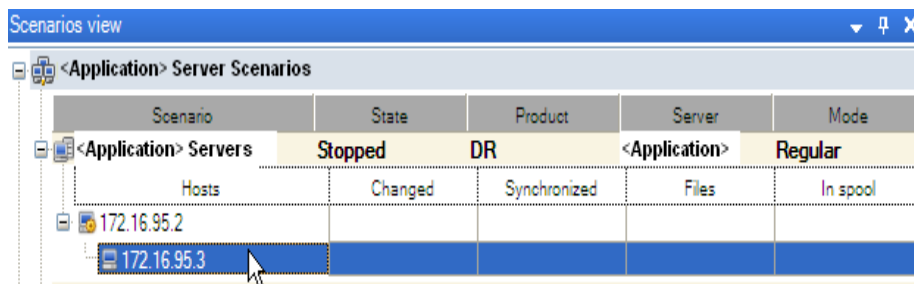
Important! The data rewind process operates in one way only - there is no replay forward. After rewind, all data subsequent to the rewind point will be lost, since data after the rewind point will be overwritten with new data.


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.

CA 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, CA 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

CA 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 34)
 - 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 34).
 - 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 36).

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 CA ARCserve RHA. AR options should be dimmed when CA 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 CA 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 CA 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 *CA-IP* in the following steps) to use Move IP redirection in your HA scenarios. This new IP address is used for CA 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.

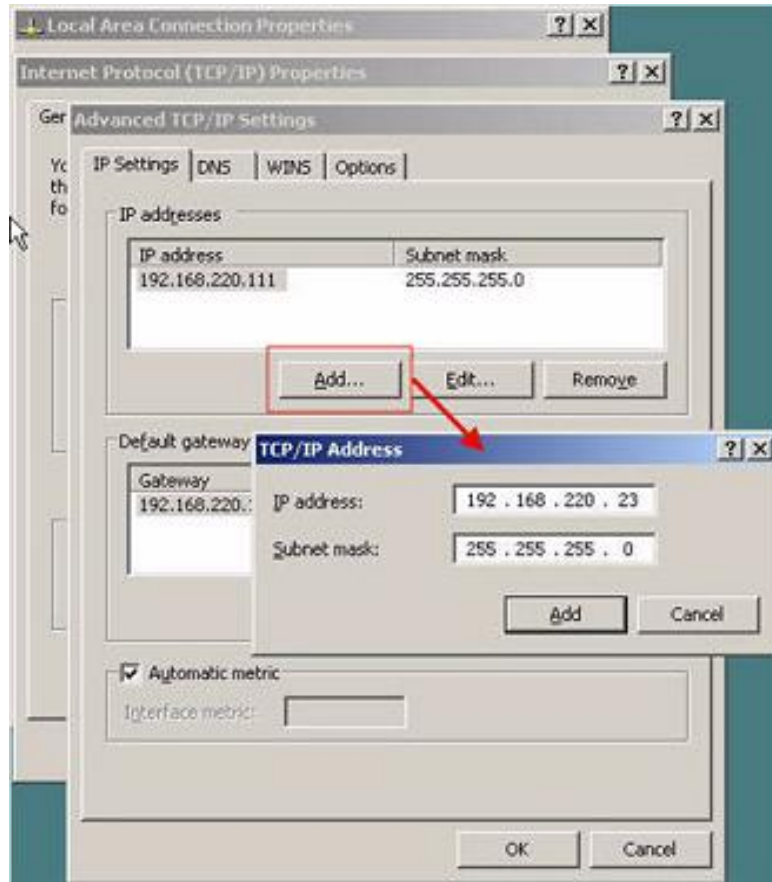
Important! Perform the following only if you are using the Move IP redirection method.

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 (CA-IP).

In the following screenshot, the CA-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 CA-IP to your HA scenarios. There are two ways to add the CA-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 CA-IP to Existing Scenarios

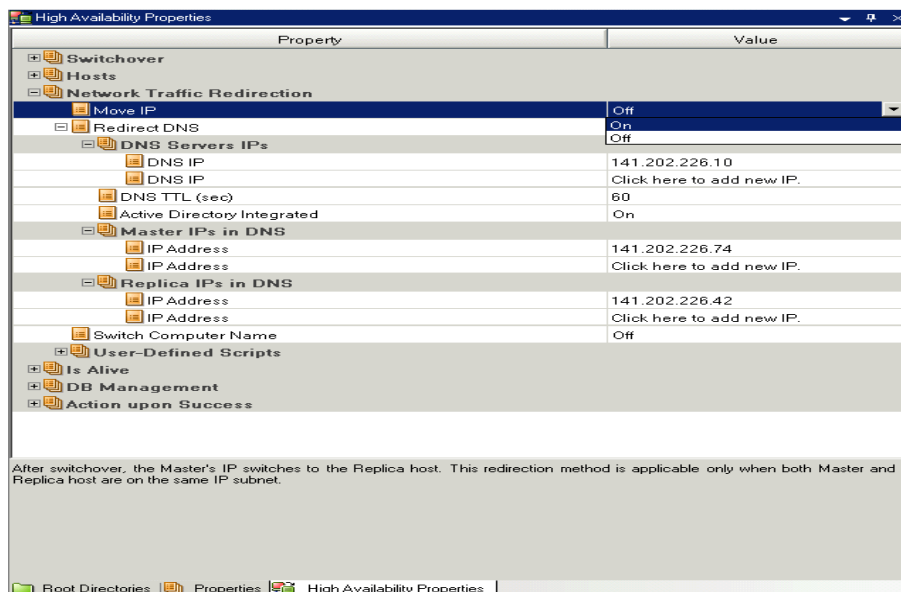
Perform this procedure only if you are using the Move IP redirection method.

To add the CA-IP to existing scenarios:

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



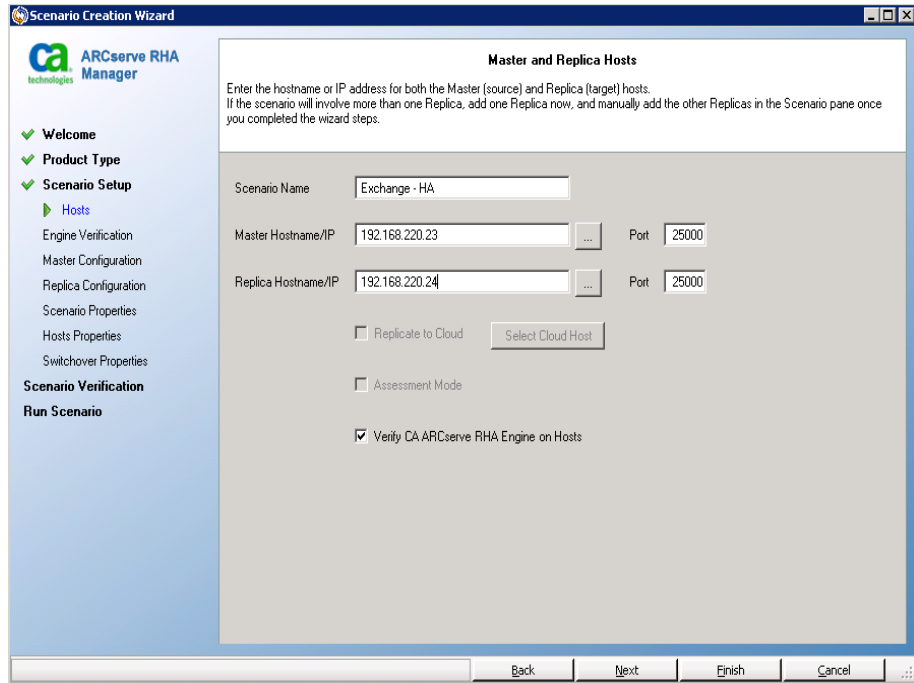
2. Right-click the Master and select Rename from the pop-up menu. Then, enter the CA-IP address.
3. On the Framework pane, select the Switchover tab and then select the Replica server as the switchover host.
4. Set the Move IP option to On. Ensure that the IP address under Move IP, IP/Mask matches the 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 CA-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 CA-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. CA 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.

CA 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, CA 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, CA 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

CA 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 *CA 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 CA 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 CA 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, CA 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 52).
- 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 *CA 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.

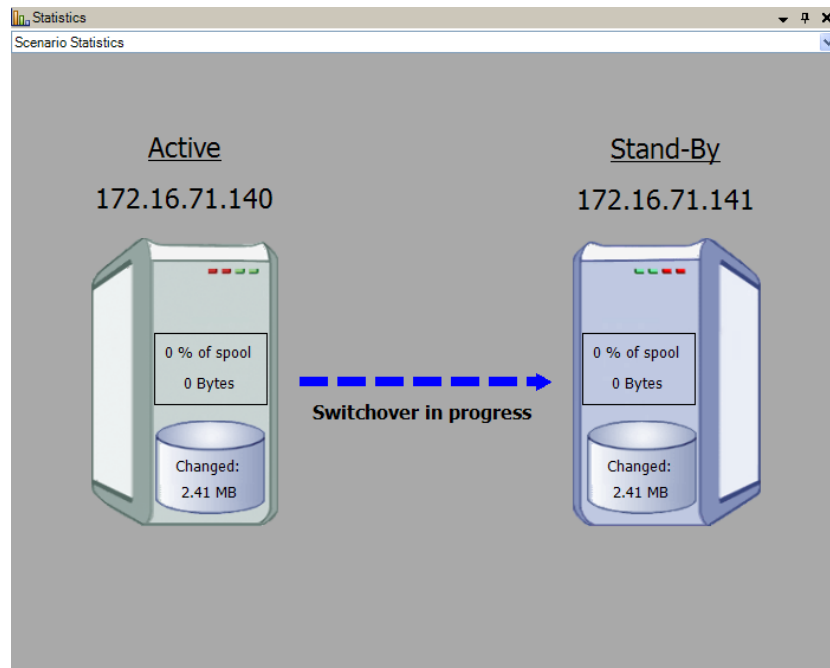
To initiate manual switchover

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.

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

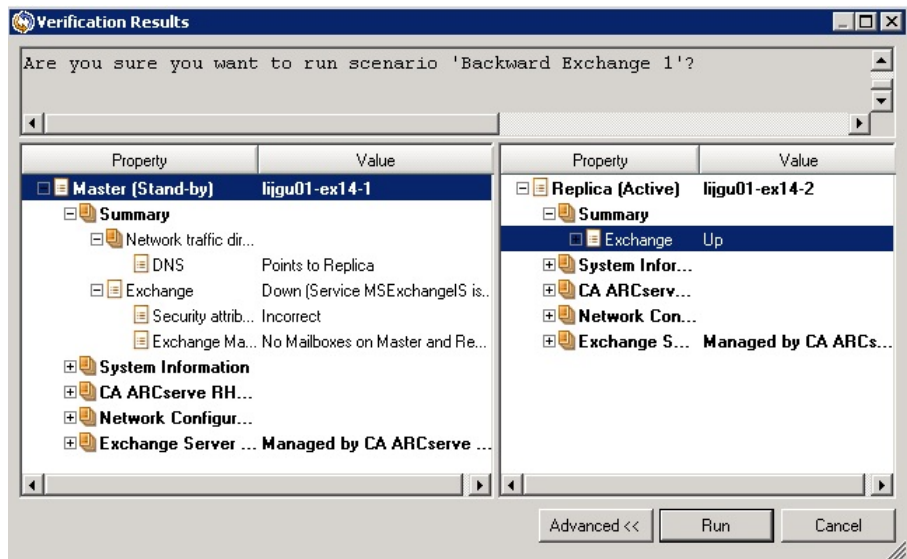
After a switchover is initiated, whether manually or automatically, at some point, you will want to reverse the server roles and make the original Master the active server and the Replica the standby server. Before you switch back the roles between servers, decide if you want the data on the original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

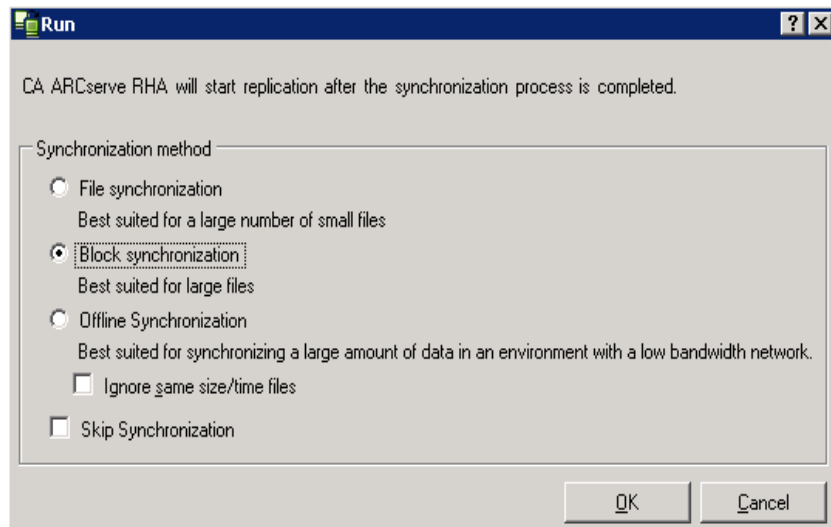
1. Ensure that both Master and Replica servers are available on the network and that the Engine is running.
2. Open the Manager and select the desired scenario from the Scenario pane.
3. Perform one of the following:
 - If the scenario is already running, skip directly to Step 4
 - If the scenario is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

CA 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.

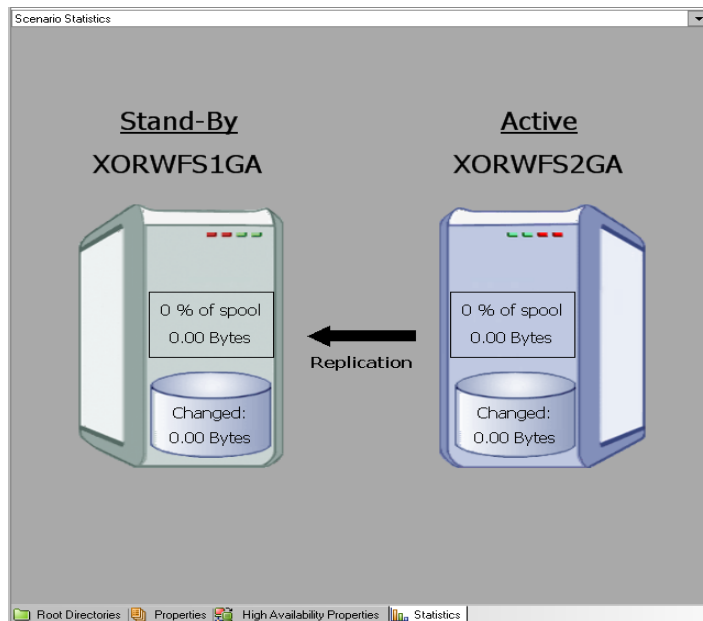


- b. Select a synchronization method from the Run dialog and click OK to start resynchronization.

Note: See the *CA 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, CA ARCserve RHA triggers Switchover without administrative involvement and could start Reverse Replication before you have investigated the cause of the failure. During Reverse Replication, CA 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.

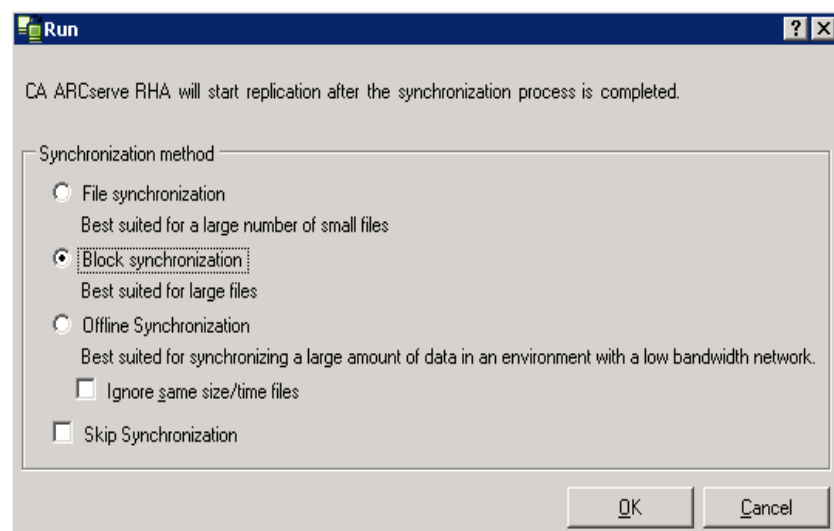
Before initiating synchronization and replication, CA ARCserve RHA verifies your scenario configuration. When verification completes successfully, the Manager displays the message: *Are you sure you want to run scenario "scenario_name?"* If problems are discovered, the top pane displays any warning and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

CA 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.

CA ARCserve RHA Report Center
Report Center Home Page
Updated: Tuesday, December 22, 2009 6:39:30 AM

Available Reports per Scenario

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

Reports
Drag a column header here to group by that column

Host	Changes	Date	Time	Type	Summary	Detailed	Size (bytes)
XORMEKCH267-1	Unknown	Today	03:29:37	Assured Recovery			811
XORMEKCH267-1	Changes found	12/07/08	22:29:48	Synchronization			29415

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)
XORMEKCH267-1	Unknown	Today	03:29:37	Assured Recovery			811
XORMEKCH267-1	Changes found	12/07/08	22:29:48	Synchronization			29415

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 screenshot displays the CA ARCserve RHA Report Center interface. At the top left is the CA Technologies logo. The main header reads "CA ARCserve RHA Report Center" and "Report Center Home Page". Below this, it says "CA ARCserve High Availability". The central heading is "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 61)

[Hyper-V Replication and High Availability](#) (see page 70)

[Switching Over and Switching Back](#) (see page 76)

[The Data Recovery Process](#) (see page 86)

[Additional Information and Tips](#) (see page 89)

[Manage Services](#) (see page 92)

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 be members of the same domain or trusted domains.
- For supported Hyper-V versions, see the CA ARCserve RHA Compatibility Matrix on CA 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 CA ARCserve RHA. You can use CA ARCserve RHA with only one Hyper-V server, but your protection will be limited to only Replication.

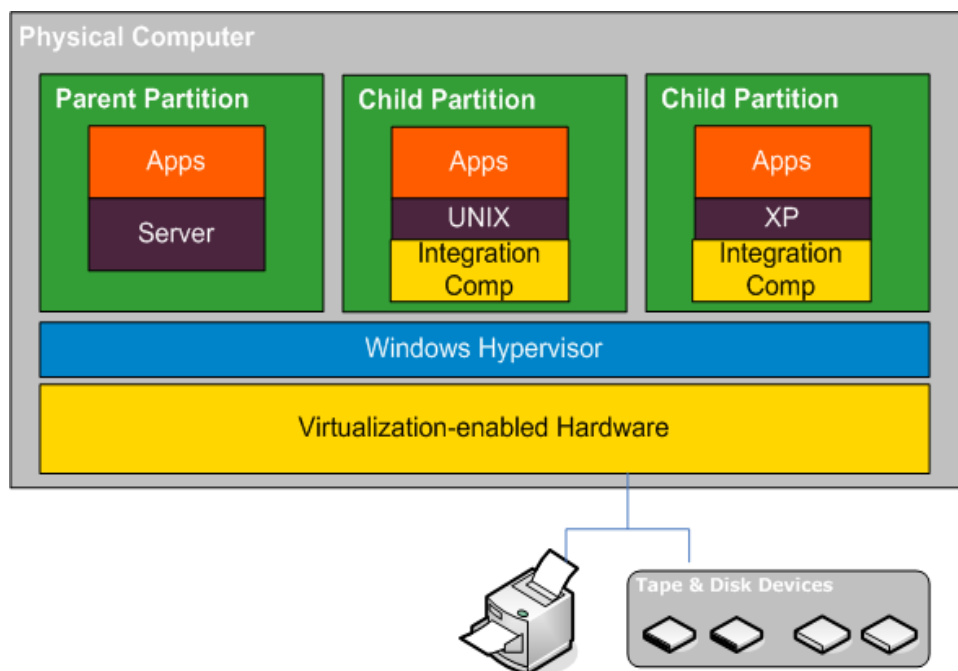
Note: For supported Hyper-V versions, see the CA ARCserve RHA Compatibility Matrix on CA Support.

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

Install CA 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. CA 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: CA 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

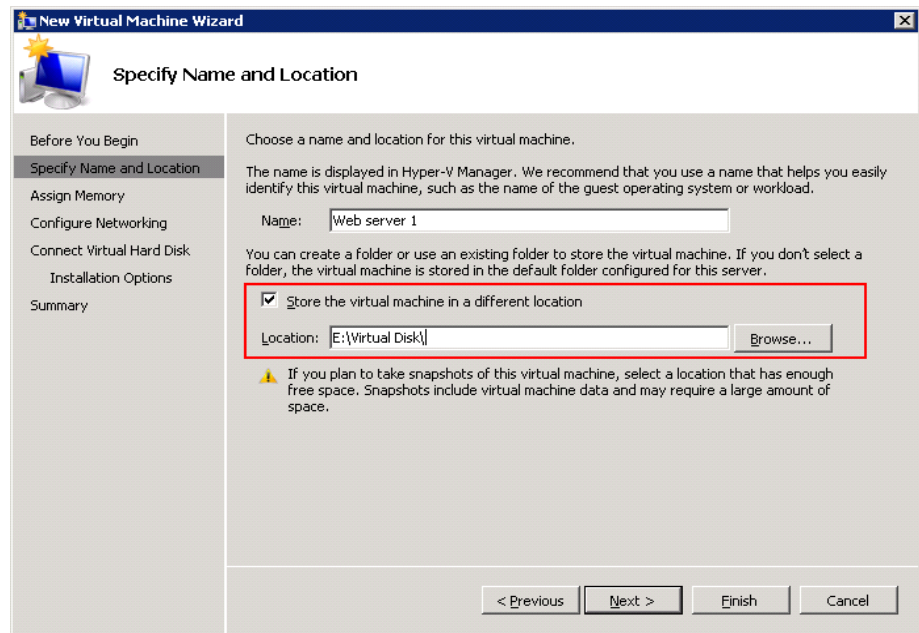
The following information applies only to CA ARCserve RHA High Availability. The default Hyper-V settings are usually sufficient when using CA ARCserve RHA but there are some default changes you should make to ensure virtual machines can easily and successfully be started on the Replica when the Master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

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

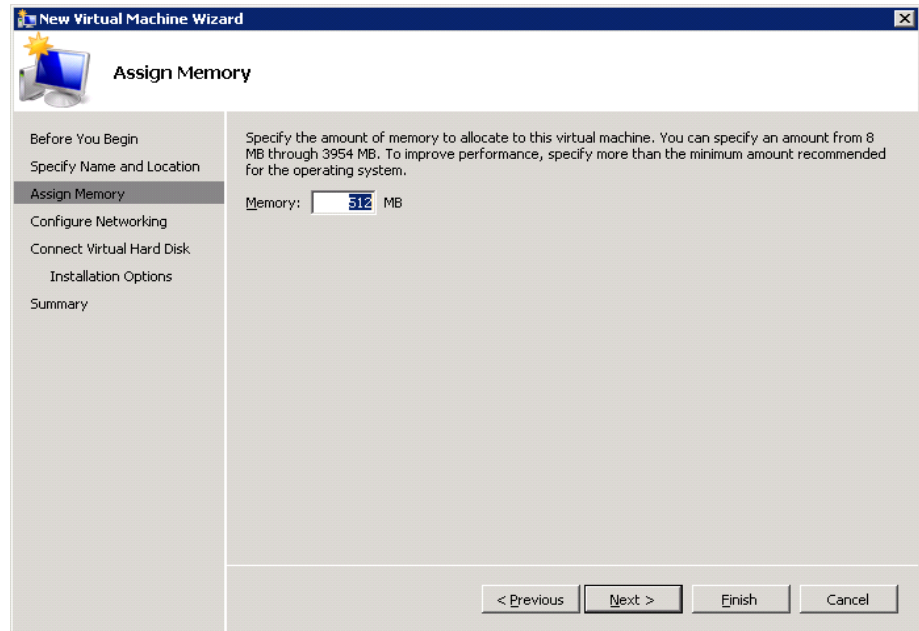
The following procedure uses a Windows guest as an example, but works for any supported Hyper-V guest operating system.

To configure Hyper-V for High Availability

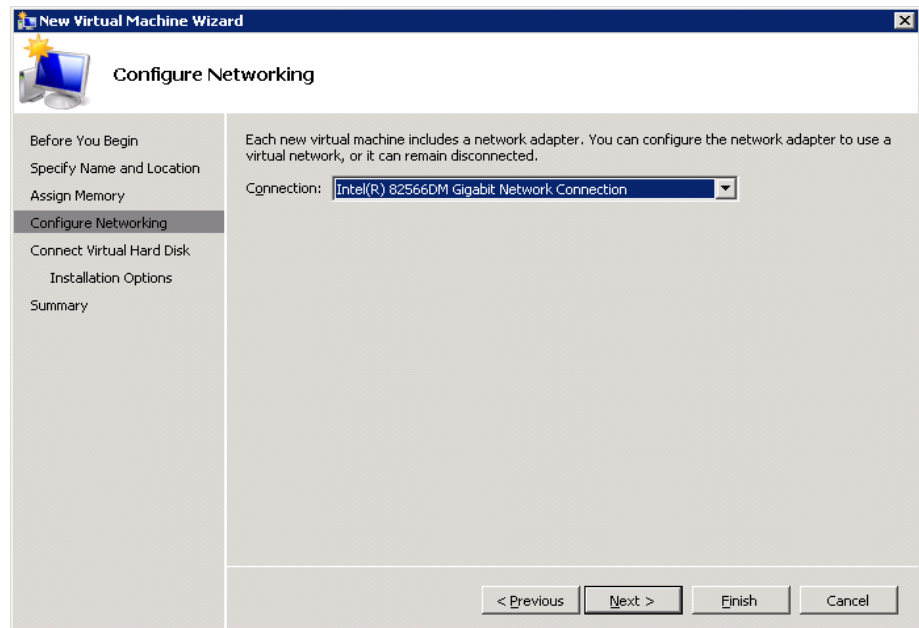
1. Launch the Hyper-V Manager wizard.
2. From the New Virtual Machine wizard dialog, enter a Name, select the **Store the virtual machine in a different location** option and then provide the location. We recommend changing the default path to a NAS or SAN volume. Click **Next**.



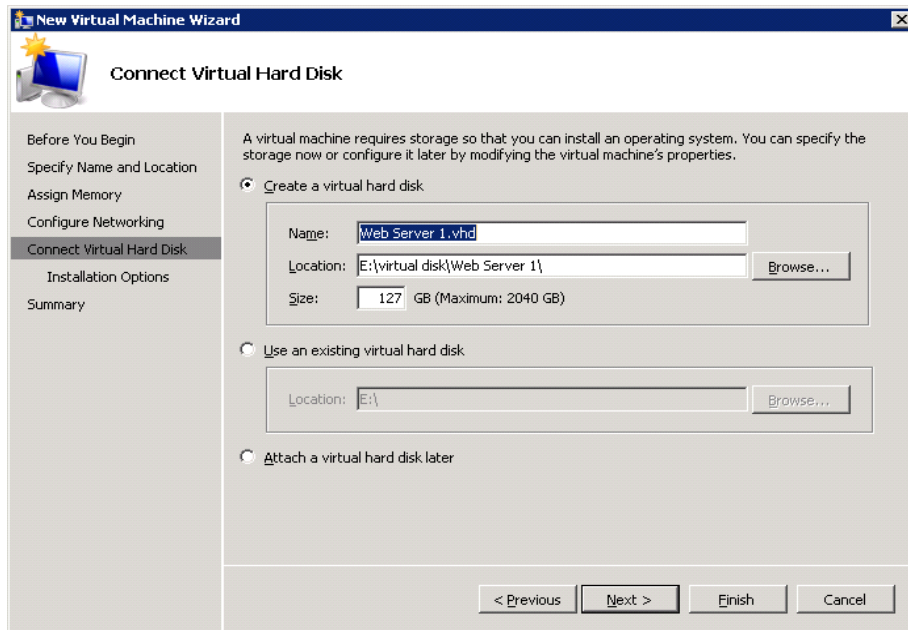
- From the Assign Memory dialog, set guest memory. Typically, 512 MB is sufficient for a single service guest. Click **Next**.



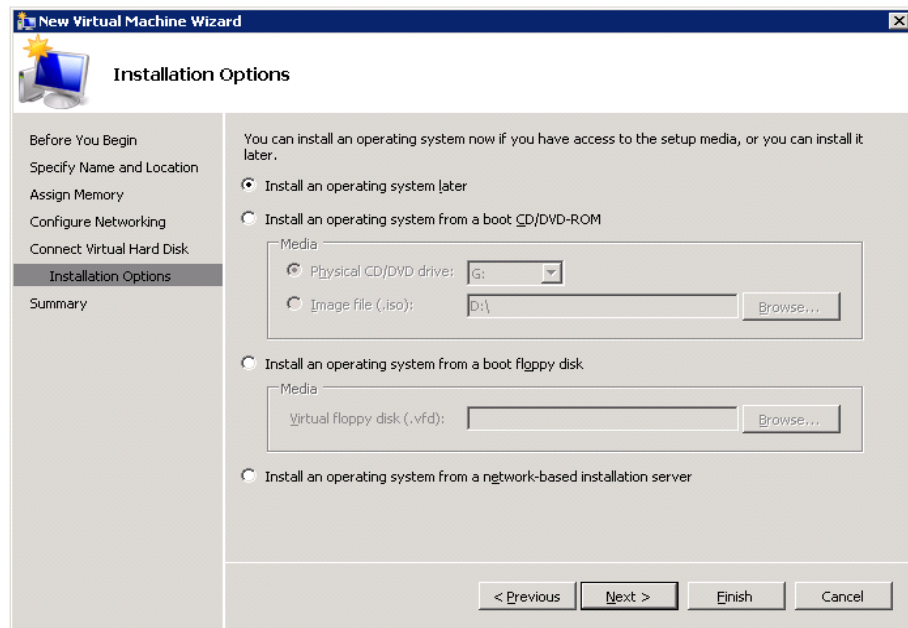
- From the Configure Networking dialog, select the virtual network adapter. Make sure you select a virtual network adapter that is bound to the physical network adapter. Click **Next**.



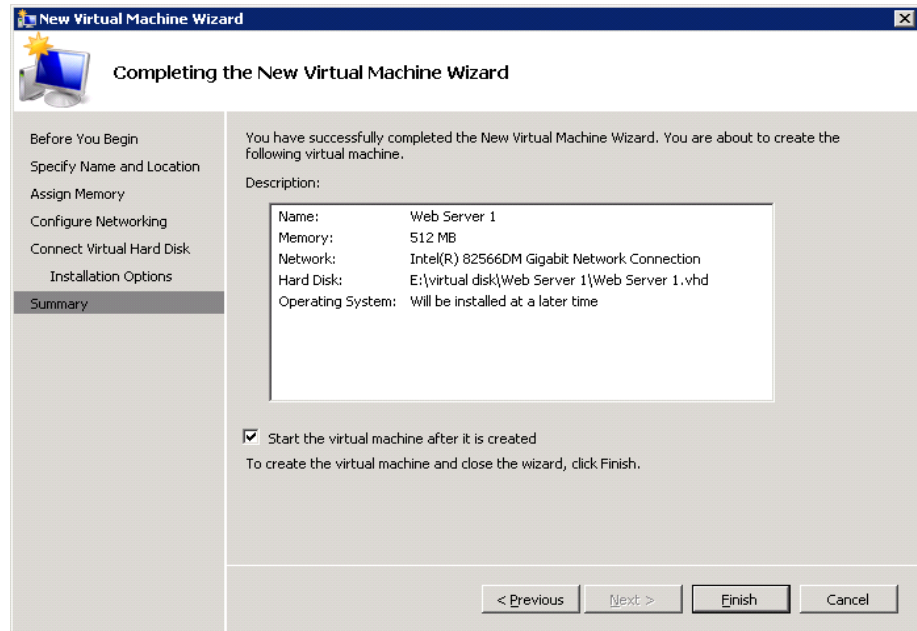
- From the Connect Virtual Hard Disk dialog, select **Create a virtual hard disk**. By default, the virtual disk is created under the VM folder. You may also use an existing disk. Click **Next**.



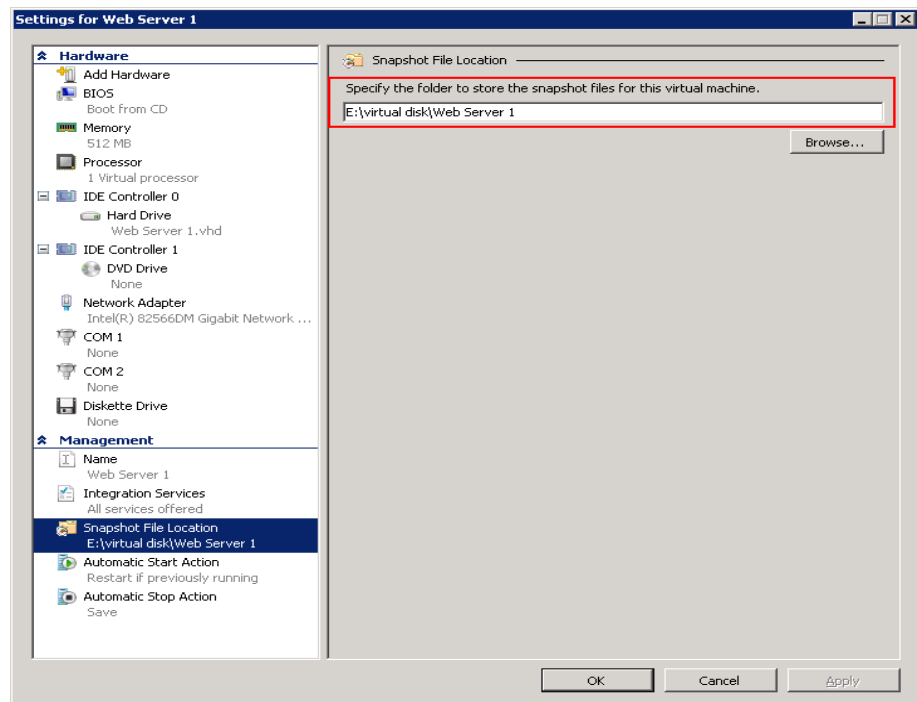
- From the Installation Options dialog, select **Install an operating system later** and click **Next**.



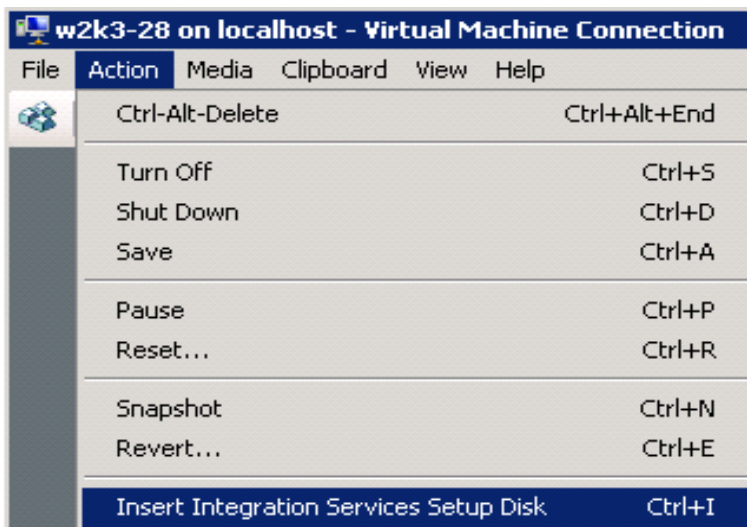
- From the Completing the New Virtual Machine Wizard dialog, select **Start the virtual machine after it is created** and click **Finish** to exit the wizard.



- 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 10)

- 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.

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 CA 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 *CA ARCserve RHA Administration Guide*.

Properties are organized into tabs on the Manager Framework pane. The tabs displayed are based upon server type, CA 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:

The screenshot displays the Hyper-V Replication Manager interface. The main window is divided into several panes:

- Scenario view:** A tree view showing the hierarchy of scenarios. The selected scenario is 'w2k3copy3', which is in a 'Running' state with 'HA' (High Availability) mode. Below it, a table lists hosts and their synchronization status.
- Scenario Statistics:** A diagram showing two servers: 'Active' (IP: 172.16.78.112) and 'Stand-By' (IP: 172.16.66.59). A blue arrow labeled 'Is Alive' points from the Stand-By server to the Active server. A black arrow labeled 'Replication' points from the Active server to the Stand-By server. Both servers show '0 % of spool' and '0.00 Bytes' of data.
- Events:** A log of events related to the scenario. The events include synchronization bookmarks, replication of modifications, synchronization reports, and synchronization completion.

Scenario	State	Product	Server	Mode
w2k3copy1	Running	DR	HyperV	Online
w2k3copy2	Running	DR	HyperV	Online
w2k3copy3	Running	HA	HyperV	Online
z832b01	Stopped ...	HA	HyperV	Online

Hosts	Changed	Synchron...	Files	In spool
155.35.78.112				
155.35.66.59				

ID	Sequence	Severity	Host/Scenario	Time	Event
IR00583	365	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/snapshots 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 *CA 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 91) 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 *CA ARCserve RHA Administration Guide*.

Properties are organized into tabs on the Manager Framework pane. The tabs displayed are based upon server type, CA 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:

The screenshot displays the Hyper-V Manager interface. The 'Scenario view' pane shows a list of scenarios:

Scenario	State	Product	Server	Mode
w2k3copy1	Running	DR	HyperV	Online
w2k3copy2	Running	DR	HyperV	Online
w2k3copy3	Running	HA	HyperV	Online
z832b01	Stopped ...	HA	HyperV	Online

The 'Statistics' pane shows a diagram of two servers: 'Active' (172.16.78.112) and 'Stand-By' (172.16.66.59). Both servers show '0 % of spool' and '0.00 Bytes' of data. A 'Replication' arrow points from the Active server to the Stand-By server, with a label 'Is Alive' above it. The Stand-By server shows 'Changed: 37.64 MB'.

The 'Events' pane at the bottom shows a list of events:

ID	Sequence	Severity	Host/Scenario	Time	Event
IR00583	365	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/snapshots 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.

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 *CA 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 91) 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 *CA 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 *CA 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. If there is only one virtual network adapter in both the Master and Replica servers, they are mapped automatically.

10. Map the virtual machines listed to the desired virtual networks on the Replica and click OK.

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 CA 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 CA 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, CA 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 78).
- 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 *CA 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.

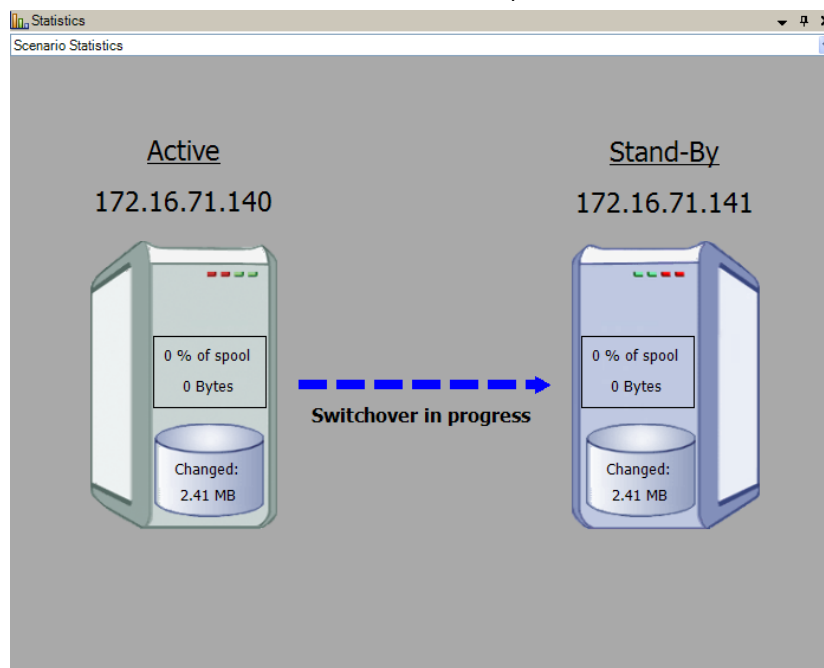
To initiate manual switchover

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

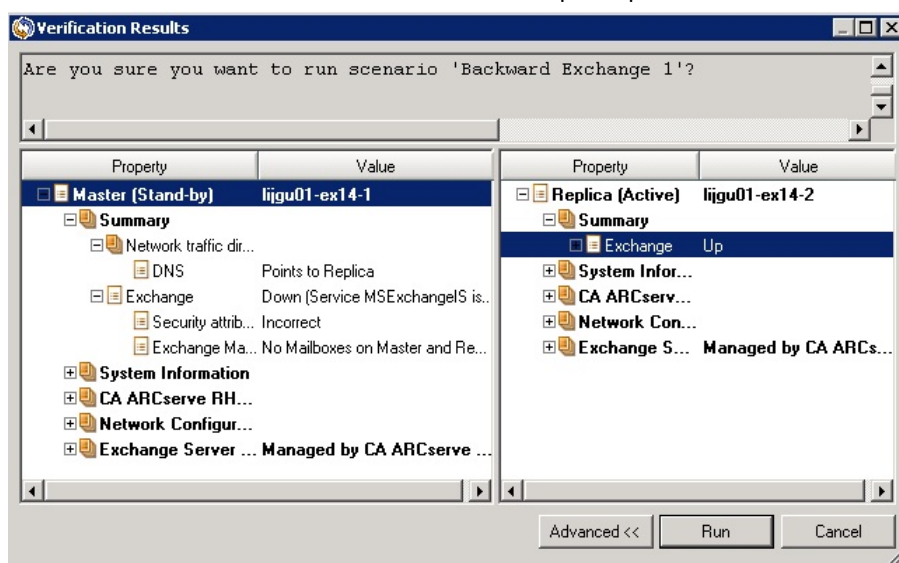
After a switchover is initiated, whether manually or automatically, at some point, you will want to reverse the server roles and make the original Master the active server and the Replica the standby server. Before you switch back the roles between servers, decide if you want the data on the original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

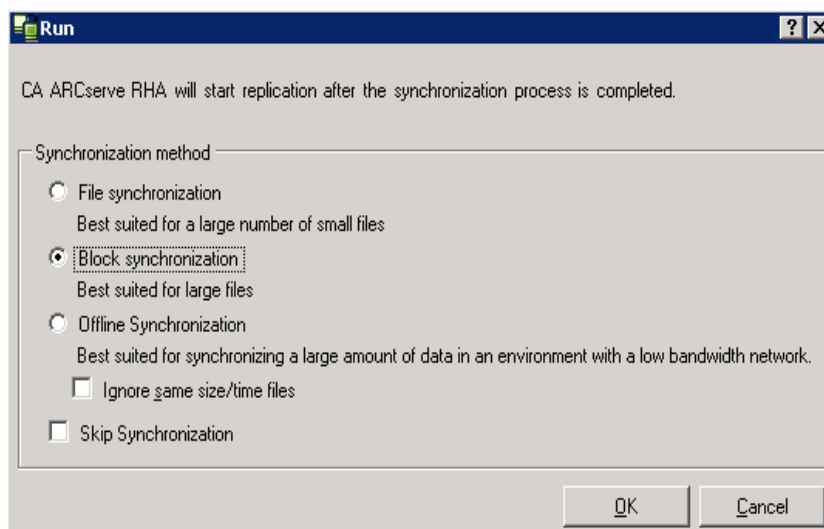
1. Ensure that both Master and Replica servers are available on the network and that the Engine is running.
2. Open the Manager and select the desired scenario from the Scenario pane.
3. Perform one of the following:
 - If the scenario is already running, skip directly to Step 4
 - If the scenario is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

CA 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.

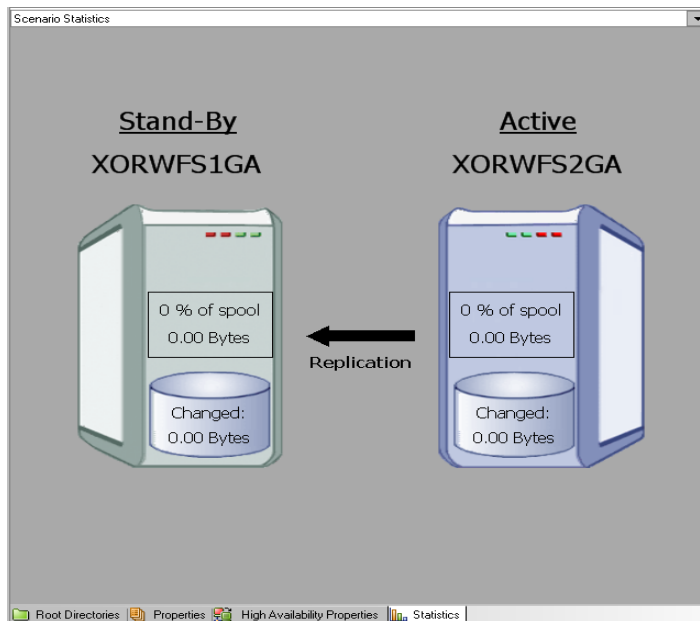


- b. Select a synchronization method from the Run dialog and click OK to start resynchronization.

Note: See the *CA 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, CA ARCserve RHA triggers Switchover without administrative involvement and could start Reverse Replication before you have investigated the cause of the failure. During Reverse Replication, CA 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.

CA 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 *CA 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 *CA 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 58).

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 *CA 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

CA 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

- 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.
- 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)
XORVMHYPERV2	Changes found	12/18/08	22:04:07	Synchronization			1923
XORVMHYPERV2	Changes found	12/18/08	04:49:54	Synchronization			2548
XORVMHYPERV2	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:

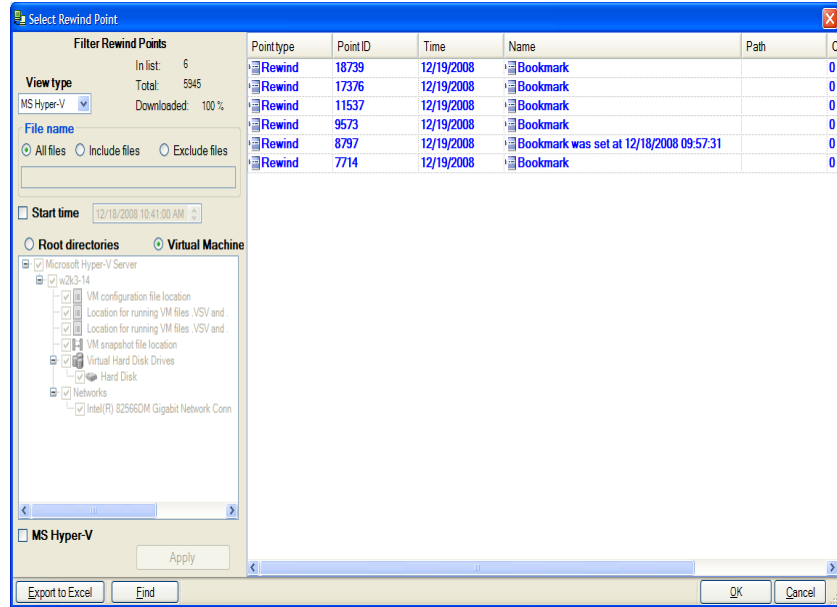
- **Stop the VM** -- CA ARCserve RHA automatically powers off the VM before recovery so that the existing VM can be overwritten. After restore completes, you need to restart the VM manually.
- **Select a bookmark** -- You must roll data back to a specific point in time called a bookmark, or rewind point. 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.

CA 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, CA 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, CA 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 CA 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.
- CA ARCserve RHA supports bandwidth limitation and bandwidth limitation scheduling. If you require such features, please consult the *CA 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 CA 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 CA 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.

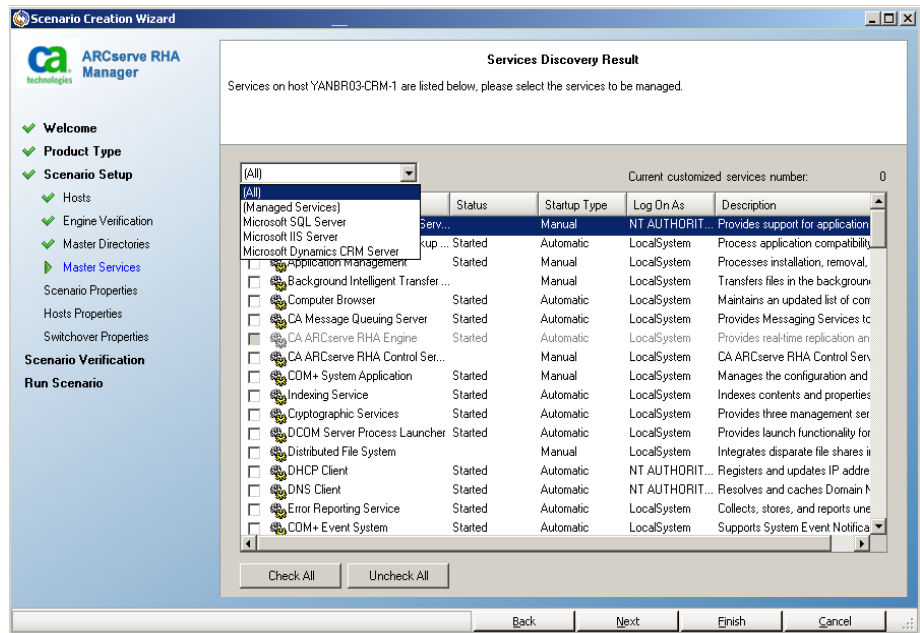
Manage Services

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 CA ARCserve RHA Manager Root Directories tab.

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

The following steps are for Custom Application scenarios.

To manage services

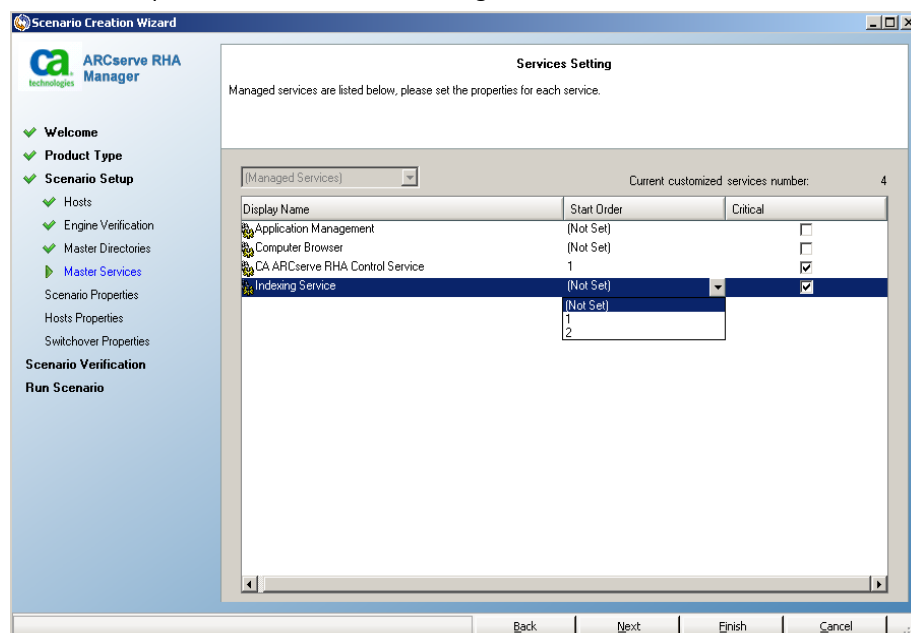


- **All** - lists all services discovered on the Master server
- **Managed Services** - lists only the checked services
- **Oracle Database** - lists Oracle-related services if the current host has Oracle installed
- **Microsoft SQL Server** - lists SQL Server-related services if the current host has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

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

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

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



3. In the Start Order column for each service you chose, specify the numeric value representing start order. For services where order does not matter, use the default value, (Not Set). The options available in the drop down list update as you configure the value. The first service has only two options: Not Set and 1. The second service has three options: Not Set, 1 and 2, and so on. If you assign the same start order to two services, CA 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.

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 95)
- [Configuring Master and Replica for Full System Scenarios](#) (see page 98)
- [Create Full System Scenarios](#) (see page 99)
- [Create Full System Scenarios for Hyper-V Platform](#) (see page 105)
- [Configure the Web Proxy to Connect to the Cloud Service](#) (see page 108)
- [Additional Properties for Full System Scenarios](#) (see page 131)
- [Configure Additional Properties in Full System-to-ESX Scenarios](#) (see page 132)
- [Configure Additional Properties in Full System EC2 Scenarios](#) (see page 133)
- [Redirection Methods for Full System Scenarios](#) (see page 135)
- [Running a Full System HA Scenario](#) (see page 135)
- [Operations on a Virtual Machine](#) (see page 136)
- [Restore Full Systems](#) (see page 138)
- [DNS Redirection using the Update DNS Tool](#) (see page 143)
- [Restoring Data to a Bare Metal Machine](#) (see page 144)
- [Restoring Data to a Bare Metal Machine after Failover](#) (see page 148)
- [Perform Full System Assured Recovery Testing](#) (see page 153)
- [Recover Active Server for Full System Scenarios](#) (see page 153)
- [Additional Tools](#) (see page 154)

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)

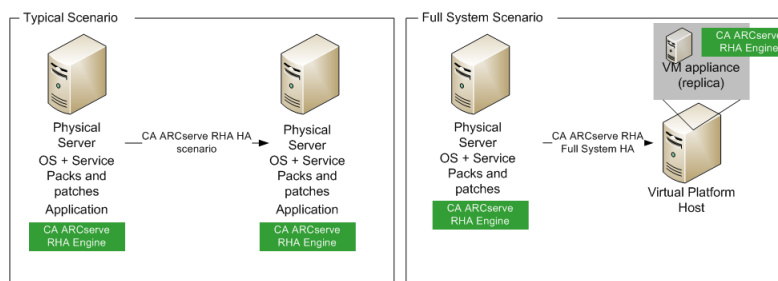
Note: Review XenServer Configuration Limitations before using XenServer on www.citrix.com <http://www.citrix.com>.

- ESX
- vCenter Server
- Amazon EC2

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

Data from a physical Master is replicated *not* to a physical Replica, as it is in typical CA ARCserve RHA scenarios, but to a virtual hard drive (VHD) 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 CA ARCserve RHA Engine. At switchover, the physical Master is disabled, a new virtual machine is created on the virtual platform host and the VHD 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 CA 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. CA 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, CA 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 ESX platforms, choose a resource pool that specifies the shared resources for the VM.
- **Datastore** -- For vCenter or ESX platforms, choose a datastore that specifies where the VM files are to be saved.

Scenario Start -- CA 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, CA ARCserve RHA initiates switchover according to the properties set during scenario creation. CA ARCserve RHA stops the P2V 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. CA 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. CA 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. CA ARCserve RHA resumes applying journal changes when AR is stopped.

Configuring Master and Replica for Full System Scenarios

Full system scenarios require three hosts instead of two, typical for other CA ARCserve RHA scenarios:

- Master server - is the host that you want to protect. This host can be physical or virtual.
- Appliance - is a VM where you installed the CA ARCserve RHA Engine.
- Virtual Platform Host - is the server where the Appliance VM is running.

To configure the Master server

For Full System scenarios, the Master can be any Windows machine, physical or virtual, supported as both a guest OS in the virtual environment and by the CA ARCserve RHA Engine. For a complete list of supported operating systems, see the CA ARCserve RHA Release Notes.

To configure the Appliance

Configure a virtual machine as follows:

- For Xen platforms, install XenServer Tools on the Appliance VM.
Note: If your virtual platform is Citrix Xen 6.0, install the .NET 4.0 Framework on the protected master. The .NET 4.0 Framework is required for installing the Xen tool on the virtual machine that the appliance created.
- For ESX and vCenter platforms, install VMware Tools on the Appliance VM.
- Install the CA ARCserve RHA Engine on the Appliance VM.
- For Amazon EC2, ensure that the date and time are correct and synchronized to your system and the network can access AWS.

The Appliance VM acts as the Replica in Full System scenarios. Data is replicated to a VHD stored in the appliance. If a failover is triggered, a new VM is created and the VHD containing replicated Master data is attached to this new VM.

Important! Do not manually delete the virtual machine, its configuration, and virtual disk files, otherwise, the scenario can fail.

Configure HTTP Communication Protocol on vCenter Server 4.0 Systems

By default, the backup proxy system and the vCenter Server systems communicate using HTTPS protocol. To specify an alternative protocol, you can configure the backup proxy system and the ESX Server system to communicate using HTTP protocol.

Note: In addition to vCenter Server 4.0 systems, the following steps apply to vCenter Server 4.1 and vCenter Server 5.0 systems.

To configure HTTP communication protocol on vCenter Server 4.0 systems

1. Log in to the vCenter Server system.

Open the file that follows using a text editor:

```
C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\proxy.xml";
```

Find the list of endpoints that contain the settings for the web service supported by the SDK.

Note: You can identify endpoints by the <EndpointList> tag.

The nested tags appear as follows:

```
<e id="5">  
<_type>vim.ProxyService.LocalServiceSpec</_type>  
<accessMode>httpsWithRedirect</accessMode>  
<port>8085</port>  
<serverNamespace>/sdk</serverNamespace>  
</e>
```

2. Change the accessMode to the following:

```
httpAndHttps
```

Close and save proxy.xml.

3. Restart the vCenter Service from the command line or from the Windows Services control panel.

Create Full System Scenarios

Full System scenarios create virtual disks from a physical Master server and store them on a virtual platform host. In the event of a switchover or failover, the virtual disk is used to create a virtual machine on the virtual platform host.

Full System scenarios are supported on the following virtual platforms:

- Hyper-V
- Citrix XenServer

Note: Review XenServer Configuration Limitations on www.citrix.com <http://www.citrix.com> before using XenServer.

- ESX
- vCenter Server
- Amazon EC2

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

While most CA ARCserve RHA scenarios require that you specify two hosts (a Master and a Replica), Full System scenarios require three:

- **Master** - this is the physical machine you wish to protect.
- **Appliance** - this is a virtual machine that acts as the Replica server (the CA ARCserve RHA Engine should be installed here). If you are using a Hyper-V virtual platform, this field does not apply and is not available (appears dim).

Note: If the Master is Windows 2008 or a later version, we recommend using Windows 2008 R2 as the appliance.

Note: If you are using Hyper-V as the destination platform, we recommend using Windows 2008 R2 as the appliance.

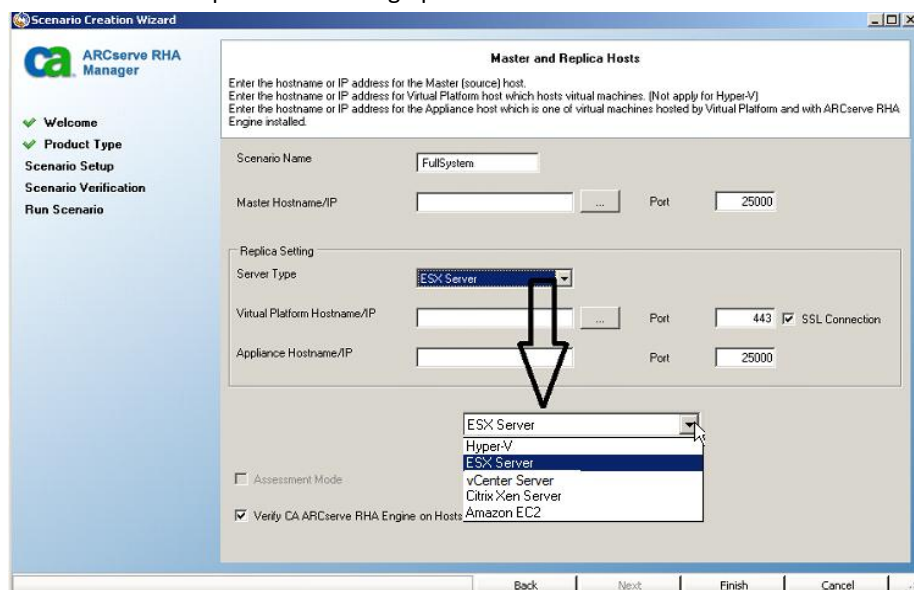
- **Virtual Platform Host** - this is the machine that hosts the Appliance VM acting as Replica server.

The following procedure applies to vCenter, ESX, and XEN Full System scenarios. For instructions on creating Hyper-V Full System scenarios, see [Create Full System Scenarios - Hyper-V](#) (see page 105).

To create full system scenarios for all platforms except Hyper-V

1. Start CA ARCserve RHA Manager. Select File, Create, New Scenario or click the New Scenario button on the toolbar.
The Welcome dialog opens.
2. Click Create a New Scenario. Type a Scenario Group name or select one from the list and then click Next.
The Select Server and Product Type dialog opens.
3. Select Full System, choose HA or DR and the desired Tasks on Replica. For more information about Tasks on Replica, see Assured Recovery. Click Next.

The Master and Replica Hosts dialog opens.



4. Complete the screen as follows and click Next when done:

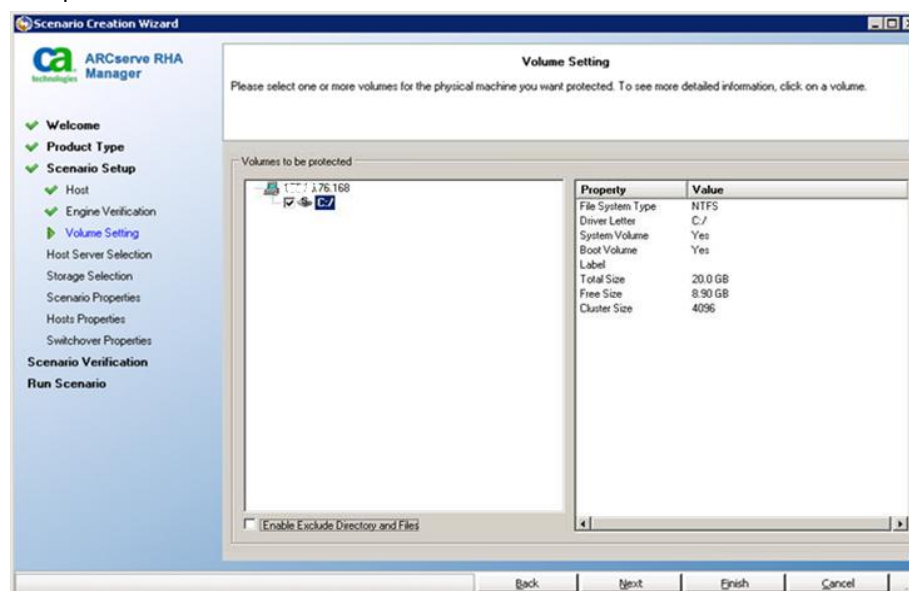
- **Scenario Name:** Type a Scenario Name. The default value is the scenario type, for example, Full System.
- **Master Hostname/IP and Port:** Specify the physical machine you wish to protect or browse to select one. Enter its port number.
- **Server Type:** Select the virtual platform of the machine that will host the VM, for example, ESX Server.
- **Virtual Platform Hostname/IP and Port:** Specify the physical machine running the virtual machine platform you selected in Server Type or browse to select one. Enter its port number.
- (Optional) **SSL Connection:** Click this option if you wish to specify an SSL port number instead. You may do so for all virtual platform types except Hyper-V.
- **Appliance Hostname/IP and Port:** Specify the virtual machine hostname or IP address of the VM to act as the Replica server in this scenario. If the server type is Hyper-V, this field is not available.
- **Verify CA ARCserve RHA Engine on Hosts:** Enable this option to confirm the latest version of the Engine is installed on all servers specified in the scenario.

Enter the appropriate credentials for the specified machines, if prompted.

5. Wait while verification completes. If desired, you may install the CA ARCserve RHA Engine on any server. If errors occur, try resolving them by clicking Verify Again. Contact your security administrator if any RPC Services errors occur.

Click Next when the Engine is verified on hosts.

The Volume Setting dialog opens. CA ARCserve RHA auto-discovers the volumes on the specified Master server.



6. Specify the volumes you wish to protect. (Optional) Enable the option, Enable Exclude Directory and Files. This option filters pagefile.sys, hiberfil.sys, System Volume Information, Recycler, and Recycled files and folders by default. Click Next.

You may be prompted to enter credentials for the server.

The Resource Pool Selection screen opens.

7. Click Next.

The Storage Selection screen opens.

8. Specify where the virtual machine should be stored. Enable the option, Allocate and commit space on demand, if desired. Click Next.

The Scenario Properties dialog opens.

9. Change properties, as desired and click Next. For more information, see the CA ARCserve RHA Administration Guide.

The Master and Replica Properties dialog opens.

10. Change properties, as desired, and click Next. The following information is acquired: CPU number, memory size on VM, adapter information on the physical Master, and the Network Mapping List. You may also change these settings outside of the scenario creation wizard. For more information, see the CA ARCserve RHA Administration Guide.
11. Wait while CA ARCserve RHA retrieves Switchover Properties.
12. When the Switchover Properties dialog opens, expand the Network Traffic Redirection properties. On this dialog, be sure to edit the physical network mappings.

To edit Physical Network Mappings:

- Select Click to edit physical network mappings.

The Network Adapter Mapping dialog opens. If there is only one virtual network adapter in both the Master and Replica servers, they are mapped automatically. Click the drop-down in the Replica Network Adapter column and choose the adapter you wish to map to the adapter listed in the Master Network Adapter column.

Apply master adapter information -- (default) Choose this option if the Master Adapter is in DHCP mode.

Customize adapter information -- Choose this option to enable the IP Settings and DNS&WINS tabs.

- **IP Settings** -- You can add or remove IP Addresses, Subnet Masks and Gateways.
- **DNS & WINS** -- You can add or remove DNS IP Addresses, Primary or Secondary WINS.

Click OK to close the Network Adapter Mappings dialog and click Next to continue.

The Switchover and Reverse Replication Initiation dialog opens.

13. Specify if switchover should be started automatically or manually. Reverse Replication cannot be specified in this scenario. Click Next.
14. 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.

15. Click Run Now if you wish to start synchronization and activate the scenario. For full system scenarios, choose Volume Synchronization. Click Finish to save current settings and run the scenario later.

Next Steps:

- For Full System-to-ESX server scenarios, you must manually enter the IP address or your scenarios will fail. For more information, see [Configure Additional Properties in Full System-to-ESX Scenarios](#) (see page 132).

Create Full System Scenarios for Hyper-V Platform

Creating Full System scenarios for Hyper-V environments is slightly different from the procedure for other virtual platforms, and displays different screens in the Scenario Creation Wizard. In Hyper-V Full System scenarios, there is no Virtual Platform Hostname or Appliance. Instead, specify the Hyper-V Hostname where the CA ARCserve RHA Engine is installed.

To create full system scenarios for Hyper-V platforms

1. Start CA ARCserve RHA Manager. Select File, Create, New Scenario or click the New Scenario button on the toolbar.

The Welcome dialog opens.

2. Click Create a New Scenario. Type a Scenario Group name or select one from the list and then click Next.

The Select Server and Product Type dialog opens.

3. Select Full System, choose High Availability and the desired Tasks on Replica. For more information about Tasks on Replica, see [Assured Recovery](#) (see page 153). Click Next.

The Master and Replica Hosts dialog opens.

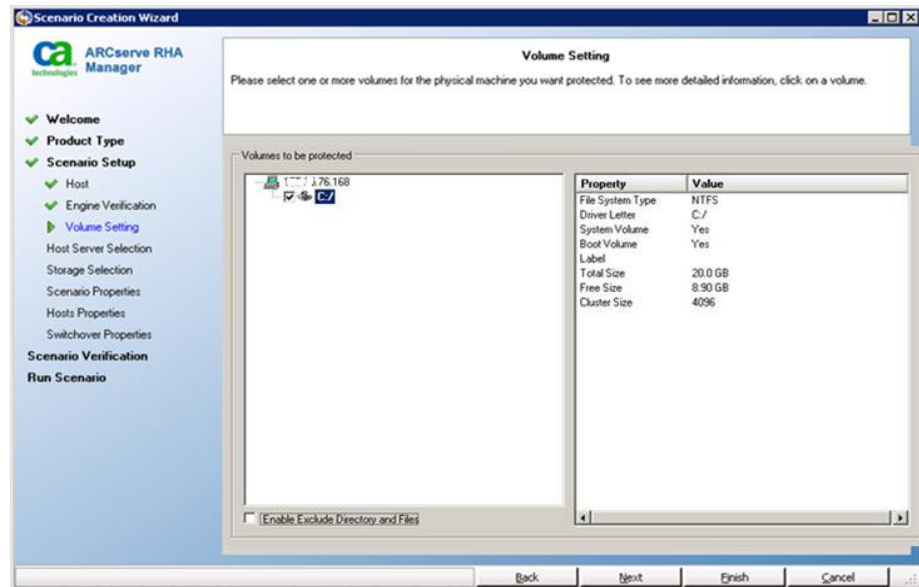
4. Complete the screen as follows and click Next when done:
 - **Scenario Name:** Type a Scenario Name. The default value is the scenario type, for example, Full System.
 - **Master Hostname/IP and Port:** Specify the physical machine you wish to protect or browse to select one. Enter its port number.

- **Server Type:** Select the virtual platform of the machine that will host the VM, for example, Hyper-V.
- **Hyper-V HostName/IP and Port:** Specify the Hyper-V host where the CA ARCserve RHA Engine is installed. This is used as the Replica server. Enter its port number.
- **Verify CA ARCserve RHA Engine on Hosts:** Enable this option to confirm the latest version of the Engine is installed on all servers specified in the scenario.

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

5. Wait while verification completes. If desired, you may install the CA ARCserve RHA Engine on the servers selected. Click Next when the Engine is verified on hosts.

The Volume Setting dialog opens. CA ARCserve RHA auto-discovers the volumes on the specified Master server.



6. Specify the volumes you wish to protect. Click Next.
7. The Storage Selection dialog opens. Specify the location where the Hyper-V VM should be stored. Enable the Allocated and commit space on demand option if you want CA ARCserve RHA to provision it. Click Next.
8. The Scenario Properties dialog opens. Change properties, as desired and then click Next. For more information, see the CA ARCserve RHA Administration Guide.

The Master and Replica Properties dialog opens.

9. Change properties, as desired and click Next. For more information, see the CA ARCserve RHA Administration Guide.

Wait while CA ARCserve RHA retrieves Switchover Properties. The following information is acquired: CPU number, memory size on VM, adapter information on the physical Master, and the Network Mapping List. You may also change these settings outside of the scenario creation wizard.

10. When the Switchover Properties dialog opens, expand the Network Traffic Redirection properties and select Click to edit physical network mapping.

The Network Adapter Mapping dialog opens. If there is only one virtual network adapter in both the Master and Replica servers, they are mapped automatically. Map the adapters listed under the Master server to the desired virtual networks on the Replica. Click the drop-down under Replica Network Adapter to select one.

- **Apply master adapter information** -- (default)
- **Customize adapter information** -- Choose this option to enable the IP Settings and DNS & WINS tabs.
 - **IP Settings** -- You may add or remove IP addresses, Subnet Masks and Gateways.
 - **DNS & WINS** -- You may add or remove DNS IP addresses, Primary and Secondary WINS.

Click OK to close the Network Adapter Mapping dialog and then click Next to continue.

The Switchover and Reverse Replication Initiation dialog opens.

11. For Hyper-V scenarios, you may only choose Switchover options. Specify to initiate switchover automatically or manually and click Next.
12. 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.

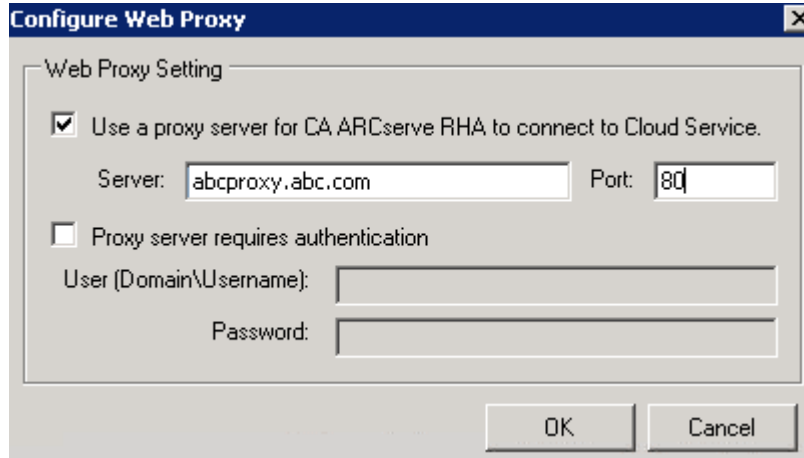
The Scenario Run dialog opens.

13. Click Run Now if you wish to start synchronization and activate the scenario. Click Finish to save current settings and run the scenario later.

Configure the Web Proxy to Connect to the Cloud Service

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


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




Add a New Cloud Account

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

To add a new cloud account

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

The Add Cloud Account dialog opens.



The screenshot shows a dialog box titled "Add Cloud Account" with the following fields and options:

- Cloud Provider: Amazon EC2 (dropdown menu)
- Cloud Account: alex@test.com (text input)
- Access Key ID: test (text input)
- Secret Access Key: [masked with asterisks] (password input)
- Use proxy to connect to the cloud service (checkbox)
- Buttons: OK, Cancel

2. Enter information in the following fields:

Cloud Provider

Specifies the name of the cloud provider service.

Cloud Account

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

Access Key ID

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


Secret Access Key

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

(Optional) Use proxy to connect to the cloud service

Lets you specify whether to use a Web proxy to communicate with the AWS web services server. If you select this check box to enable this option, ensure that you have first [configured](#) (see page 108) this proxy.

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



Access Credentials

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

Access Keys | X.509 Certificates | Key Pairs

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

Your Access Keys

Created	Access Key ID	Secret Access Key	Status
January 27, 2009	003KNR20D32SJNAQ5ET2	Show	Active (Make Inactive)
October 11, 2010	A5ETY8A6DJS2A92NSKA6	Show	Active (Make Inactive)

[View Your Deleted Access Keys](#)


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

[Learn more about Access Keys](#)


3. Click OK.

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


Change the Default AWS Region

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

Update Cloud Account Information

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

Delete a Cloud Account

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

Create a New EC2 Replica Instance

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

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


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

- The AWS account is registered in the CA ARCserve RHA Manager

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

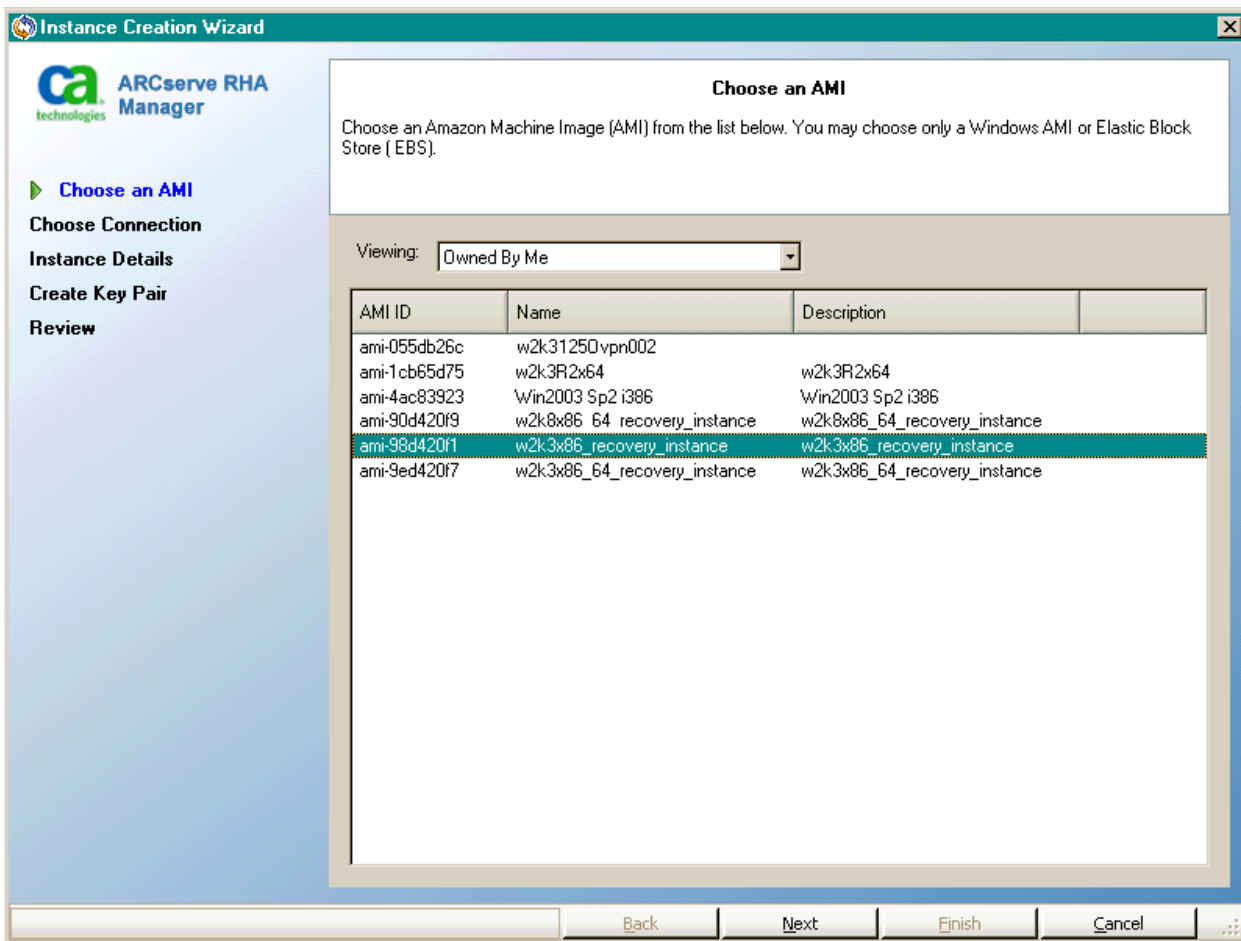
To create a new EC2 replica instance

1. Select the Cloud View panel in the CA ARCserve RHA Manager.

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

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

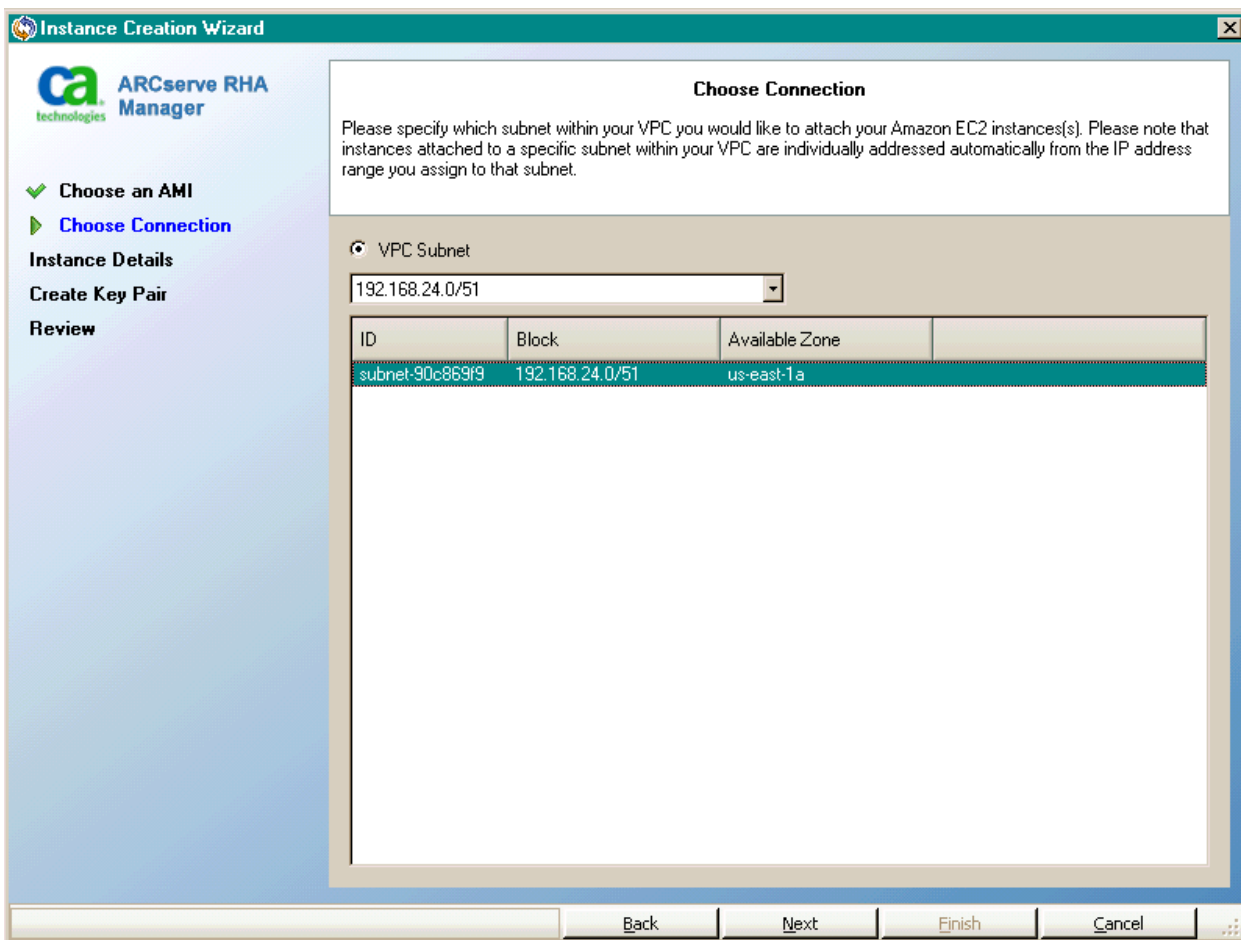
The Instance Creation Wizard opens.



3. Select an Amazon Machine Image (AMI) from the list in the Choose an AMI dialog and click Next.

Note: You can only use EBS-backed Windows AMIs.


The Choose Connection screen opens.



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

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

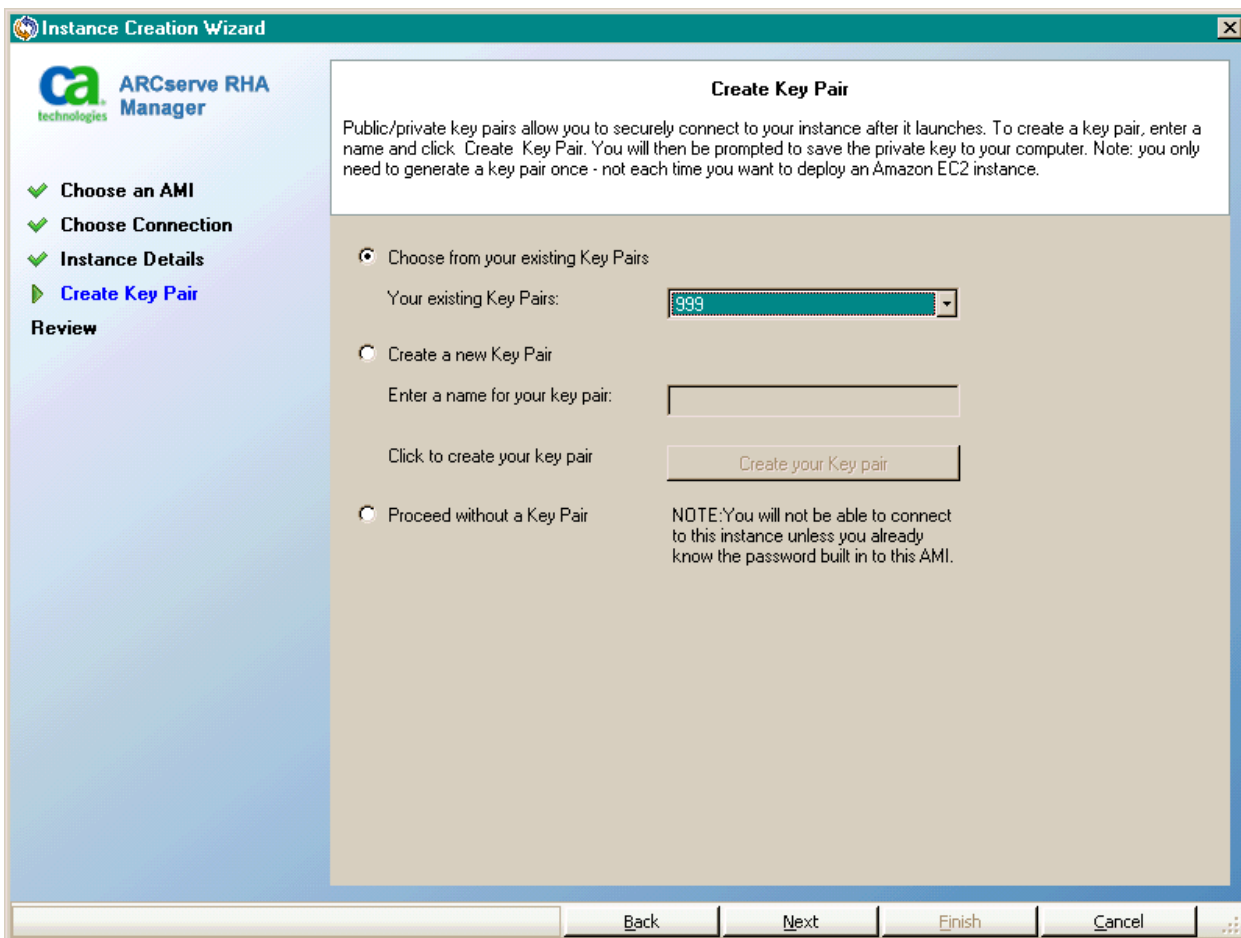
The Instance Details screen opens.



The screenshot shows the 'Instance Creation Wizard' window with the 'Instance Details' step selected. The window title is 'Instance Creation Wizard' and the logo for 'ARCserve RHA Manager' is visible in the top left. The left sidebar contains a navigation menu with the following items: 'Choose an AMI' (checked), 'Choose Connection' (checked), 'Instance Details' (selected), 'Create Key Pair', and 'Review'. The main content area is titled 'Instance Details' and contains the following text: 'Provide the details for your instance(s). EC2 Instances let you pay for compute capacity by the hour with no long term commitments. This transforms what are commonly large fixed costs into much smaller variable costs.' Below this text are four input fields: 'Number of Instances' (text box with '1'), 'Availability Zone' (dropdown menu with 'us-east-1a'), 'Instance Type' (dropdown menu with 'Small (m1.small 1.7 GB)'), and 'Instance Name' (text box with 'test-01'). At the bottom of the window are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

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

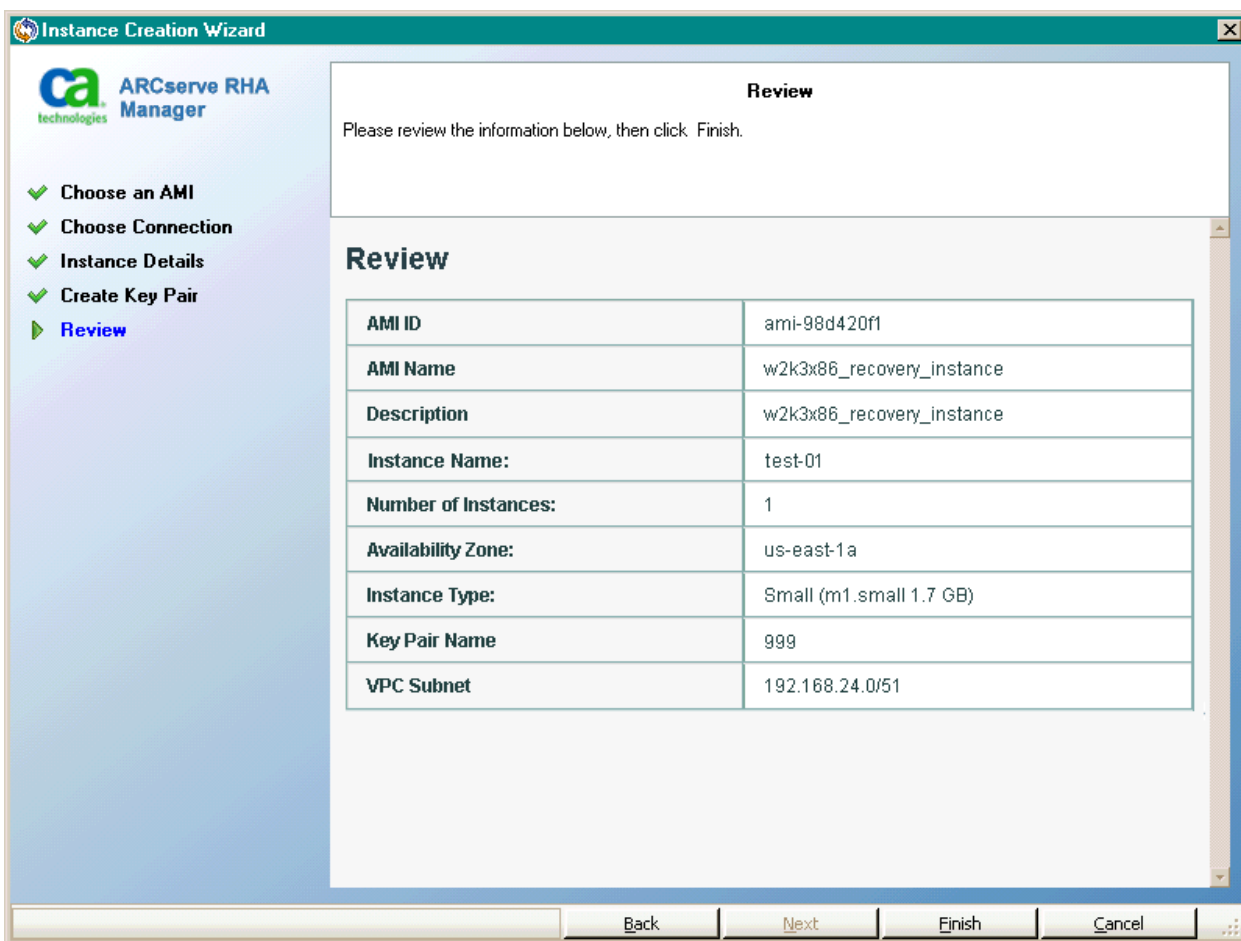
The Create Key Pair screen opens.



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

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

The Review screen opens.




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

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


Notes:

- The Cloud View panel only allows creation of VPC-based instances. You cannot allocate an instance outside VPC, so VPC configuration is mandatory.
- You can create non-VPC instances (public instances) using the AWS management tools (in the AWS Management Console). These public instances will be displayed in the Cloud View panel. However, those instances will not be available in the Scenario Creation Wizard because the wizard only supports “private to private” scenarios, or replication from a private on-premises network to VPC.
- You can filter resource browsing by selecting different regions. There are seven regions available to AWS users: EU West (Ireland), US East (Virginia), US West (N. California), US west (Oregon), South america (Sao Paulo), Asia Pacific (Tokyo), and Asia Pacific (Singapore). Currently AWS only allows creation of one VPC per AWS account. Every VPC may have several subnetworks associated with it (you select the subnetwork ID when the instance is allocated). The subnetwork may reside in one of the four regions. If you want to allocate an instance in a specific subnetwork, then you should first select the corresponding region in the *Change Default Region* drop-down, however, note that RHA only supports two regions: US East (Virginia) and EU West (Ireland).


Start an EC2 Replica Instance

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

Stop an EC2 Replica Instance

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

Delete an EC2 Replica Instance

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

Create a New EC2 Data Replication Scenario

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

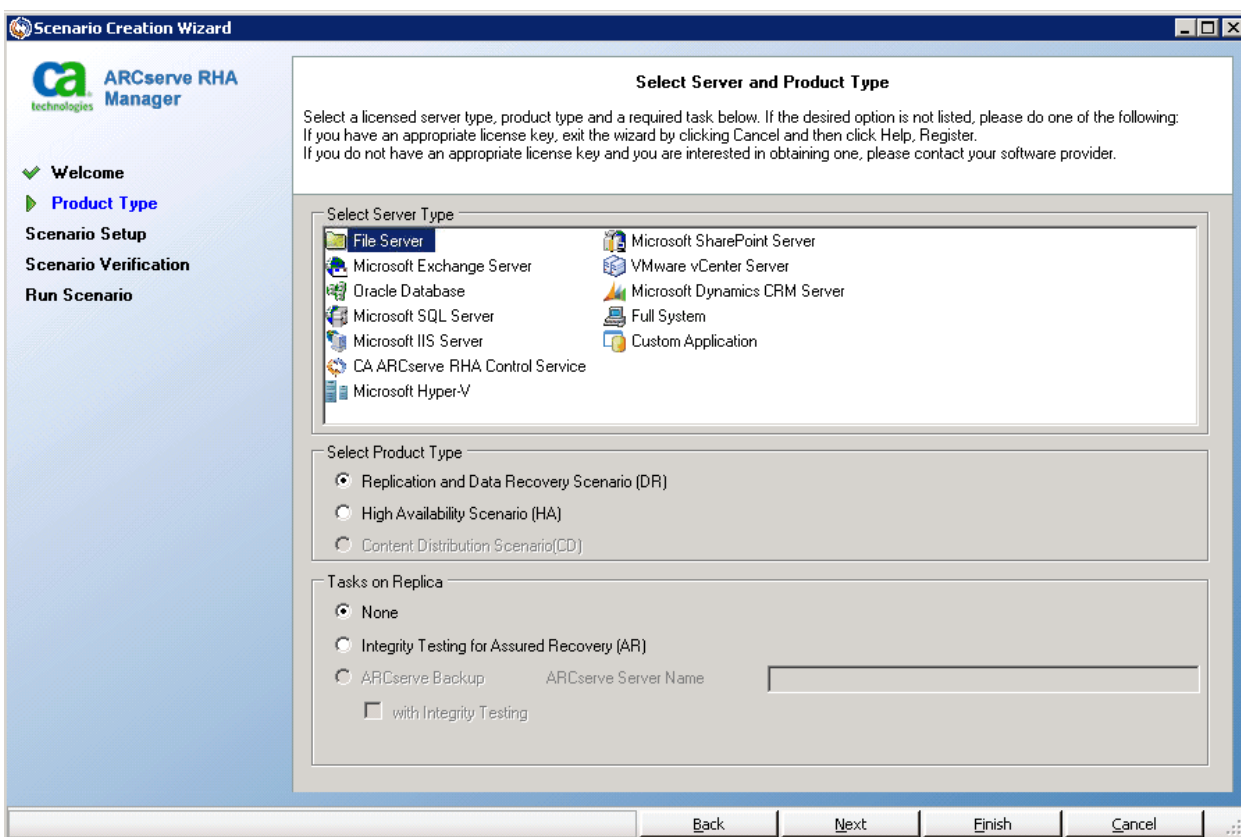
To create a new EC2 Data Replication scenario

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

The Welcome screen opens.

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

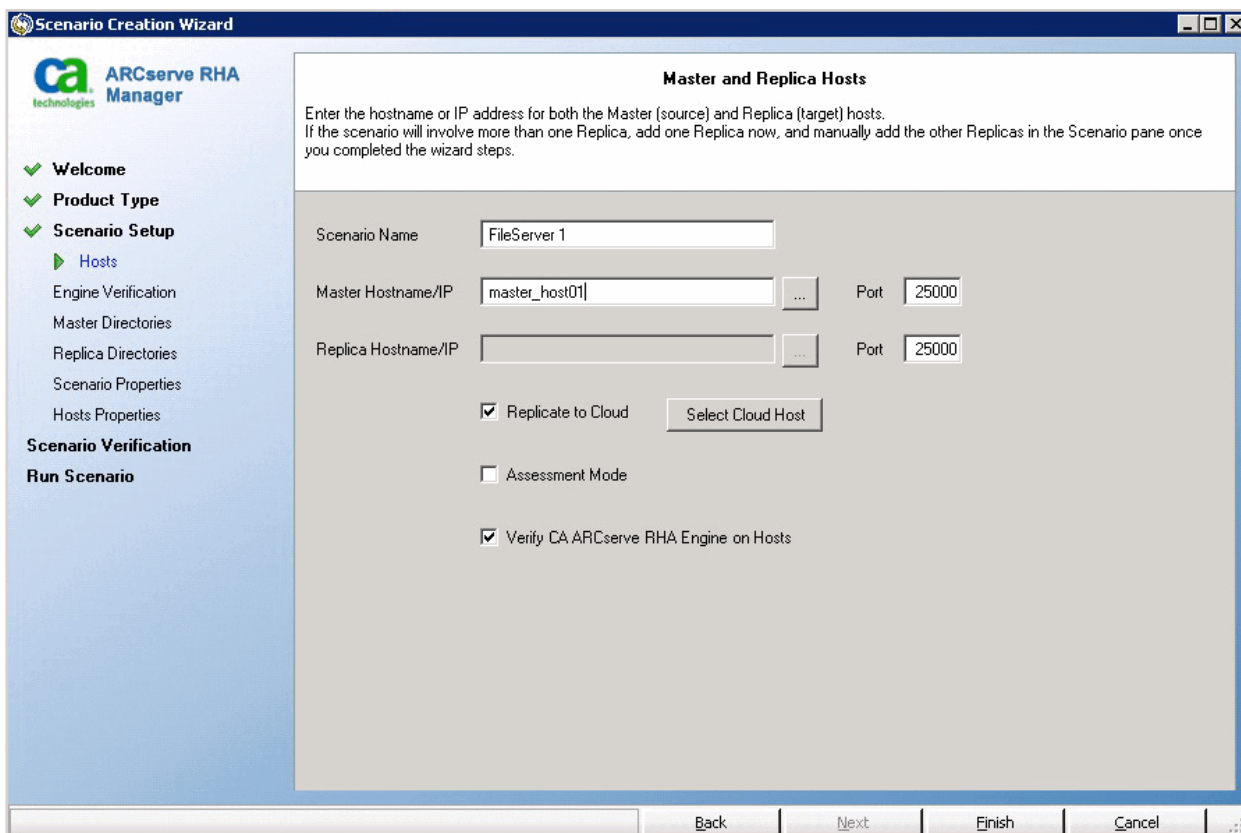
The Select Server and Product Type screen opens.



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

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

The Master and Replica Hosts screen opens.




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

The Cloud Instance Selection dialog opens.

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

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

The Engine Verification screen opens.



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

The Volume Setting screen opens.

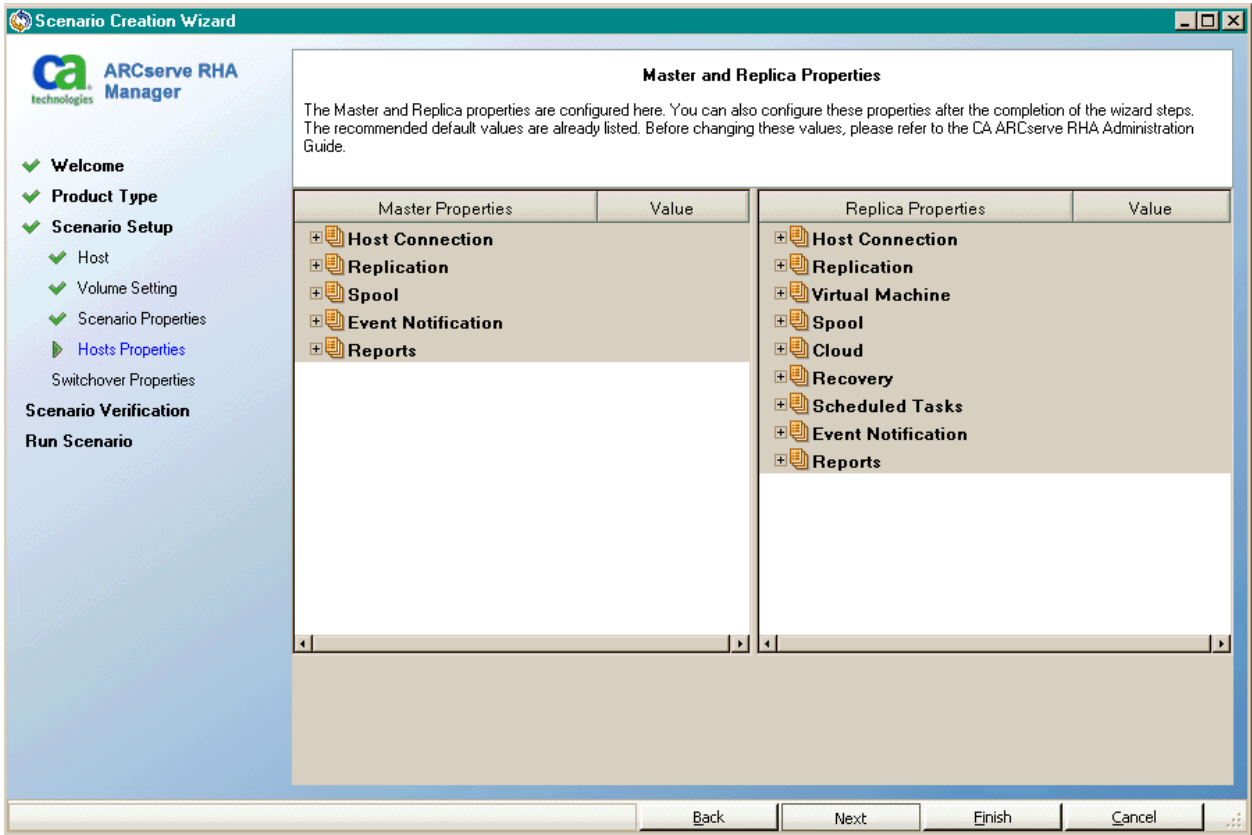
7. Enter information and click Next.

The Scenario Properties screen opens.

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

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

The Master and Replica Properties screen opens.



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

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

10. Click Next.
Wait for Scenario Verification to complete.
11. If errors or warnings are listed, resolve them before continuing. When ready, click Next.
The Scenario Run screen opens.
12. Click Run Now to start synchronization and activate the scenario, or click Finish to run the scenario later.

Create a New Full System EC2 High Availability Scenario

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

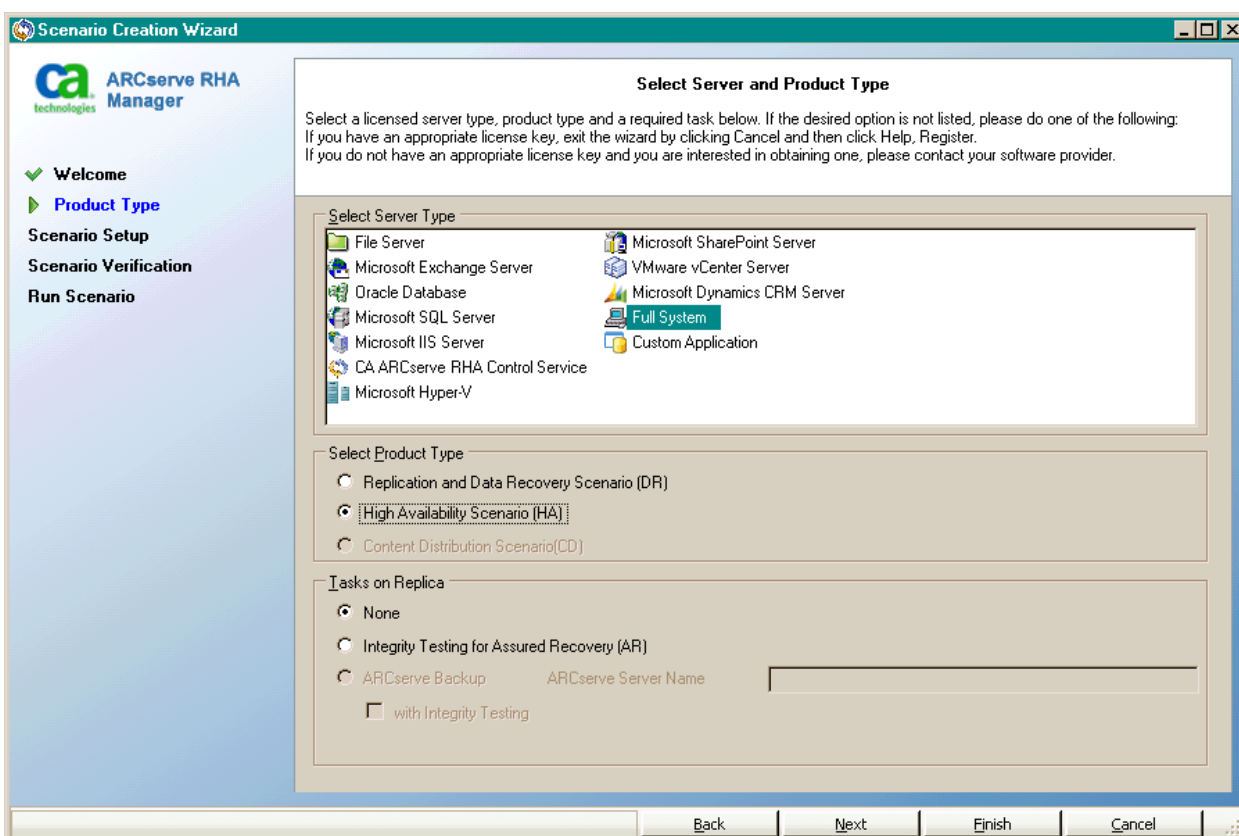
To create a new full system EC2 High Availability Scenario

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

The Welcome screen opens.

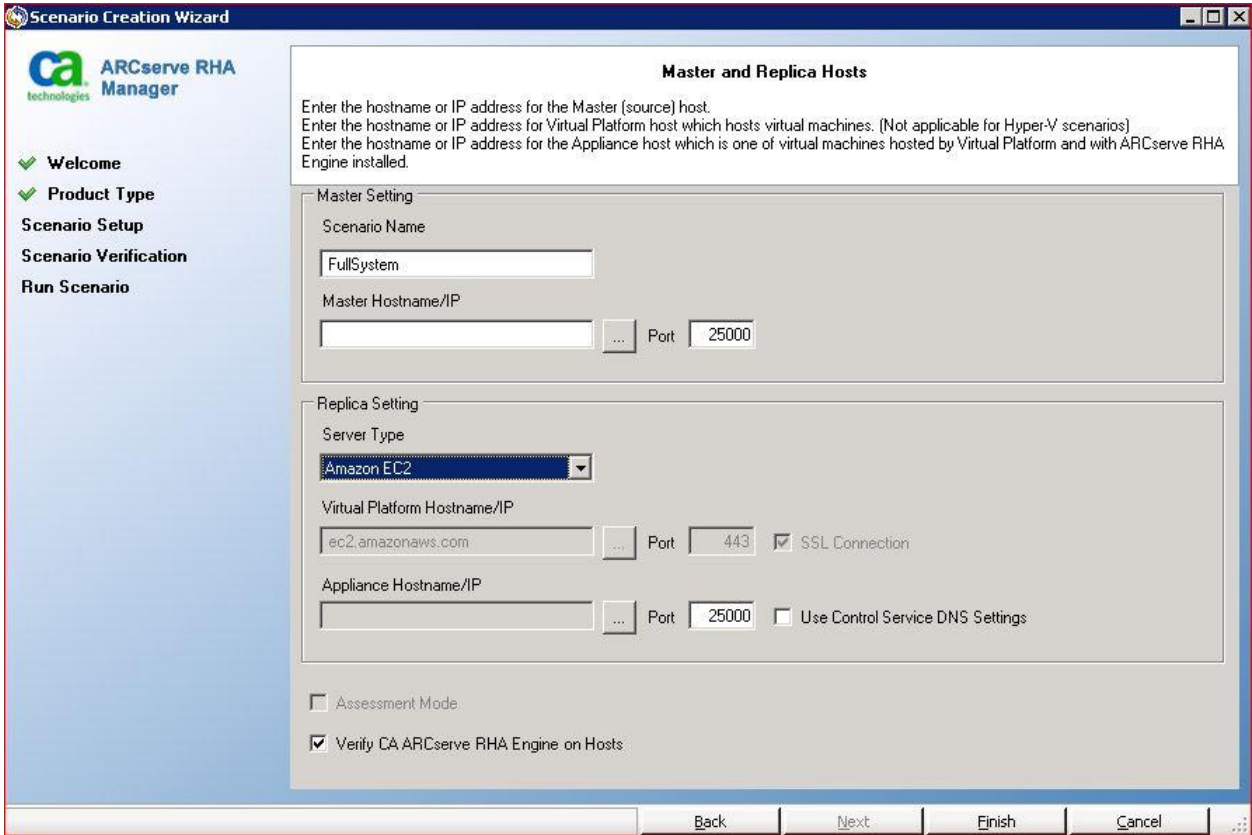
2. Choose Create a New Scenario, select a Group from the list and then click Next.

The Select Server and Product Type screen opens.



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

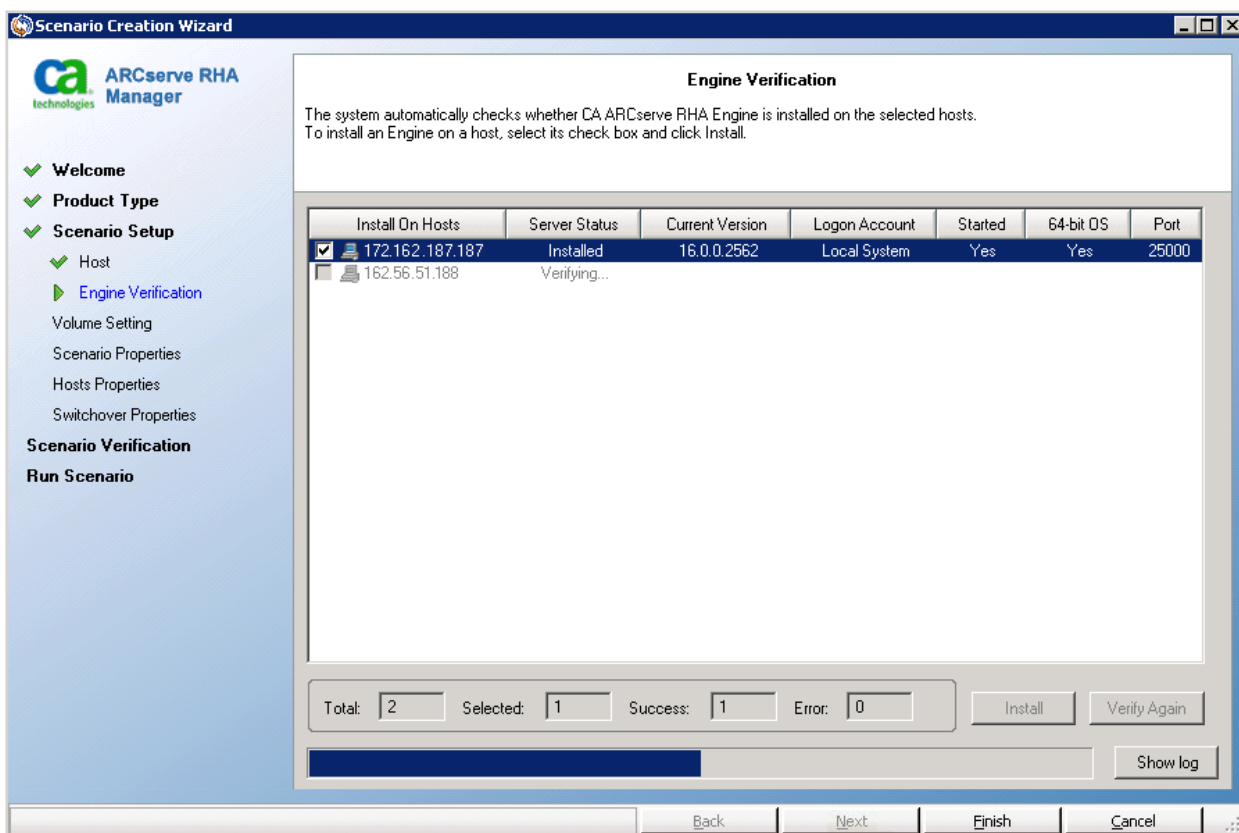
The Master and Replica Hosts screen opens.



4. Do the following in the Master and Replica Hosts screen:
 - a. Type a Scenario Name and enter the Hostname or IP Address and Port number for the Master server.
 - b. Specify Amazon EC2 as the Replica server.
 - c. Specify the EC2 replica instance (appliance). Click the button to browse for and select the AWS account and EC2 replica instance (appliance).

The Cloud Instance Selection dialog opens.
 - d. Select the AWS account, cloud replica (appliance), and region and click OK.
 - e. Select or clear the Use Control Service DNS check box as needed. Selecting this check box applies the DNS settings from the Control Service server to the EC2 replica instance (appliance) host.
 - f. Ensure the Verify CA ARCserve RHA Engine on Hosts option is enabled (default) and click Next.

The Engine Verification screen opens.



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

The Volume Setting screen opens.

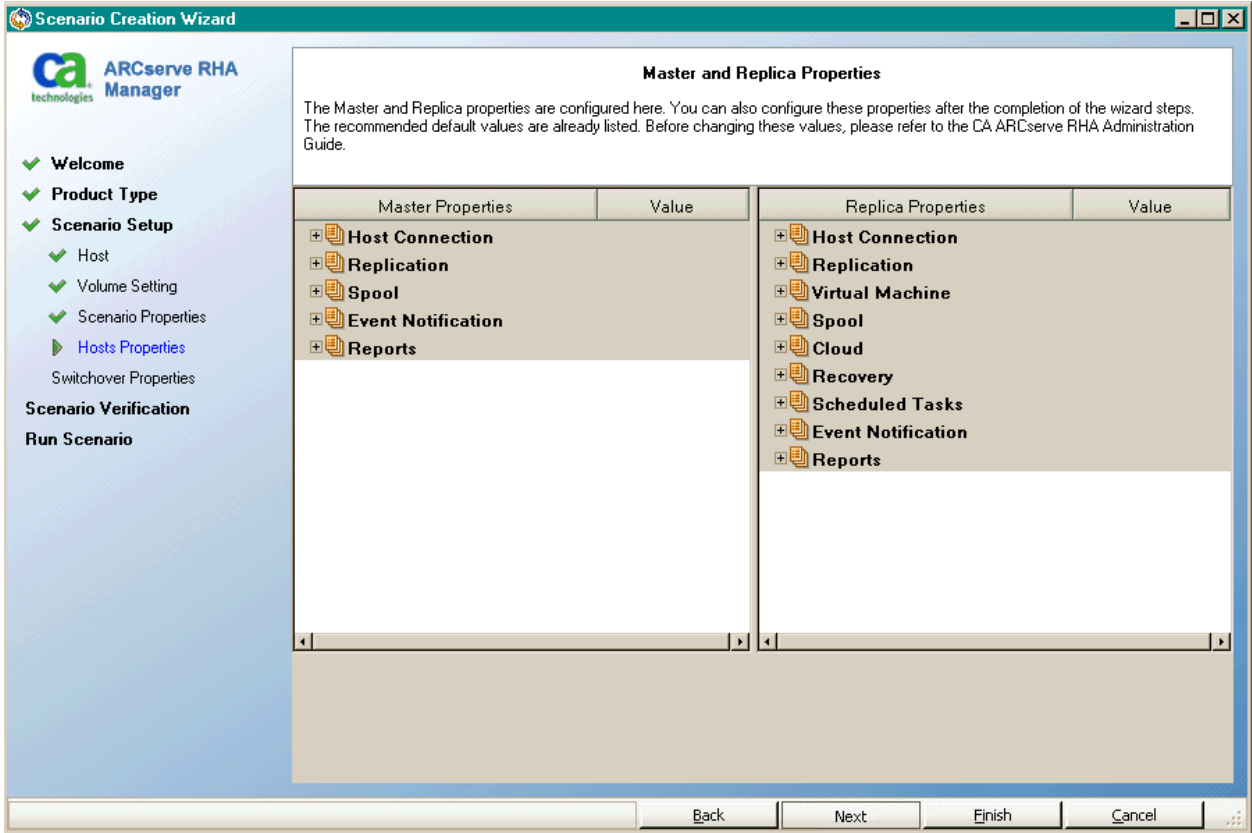
6. Select one or more volumes for the physical machine you want to protect and click Next.

The Scenario Properties screen opens.

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

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

The Master and Replica Properties screen opens.



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

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

The Switchover Properties screen opens.

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

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

Switchover Properties

Switchover properties will be configured in this step. The recommended default values are listed below. Refer to the Administration Guide before modifying the current settings.

Property	Value
Switchover	
Switchover Hostname	197.178.11.49
Hosts	
Master Fully Qualified Name	test01-sql.local
Replica Fully Qualified Name	ip-C0A83114
Network Traffic Redirection	
Redirect DNS	On
DNS Servers IPs	
DNS IP	162.14.2.191
DNS IP	162.14.1.76
DNS IP	Click here to add new IP.
DNS TTL (sec)	60
Active Directory Integrated	On
Master IPs in DNS	
IP Address	162.46.2.106

Enter IPs of DNS servers to update. CA ARCserve High Availability tries to update all servers listed. However, switchover is considered successful even if only one update is successful.

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

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

Switchover Properties

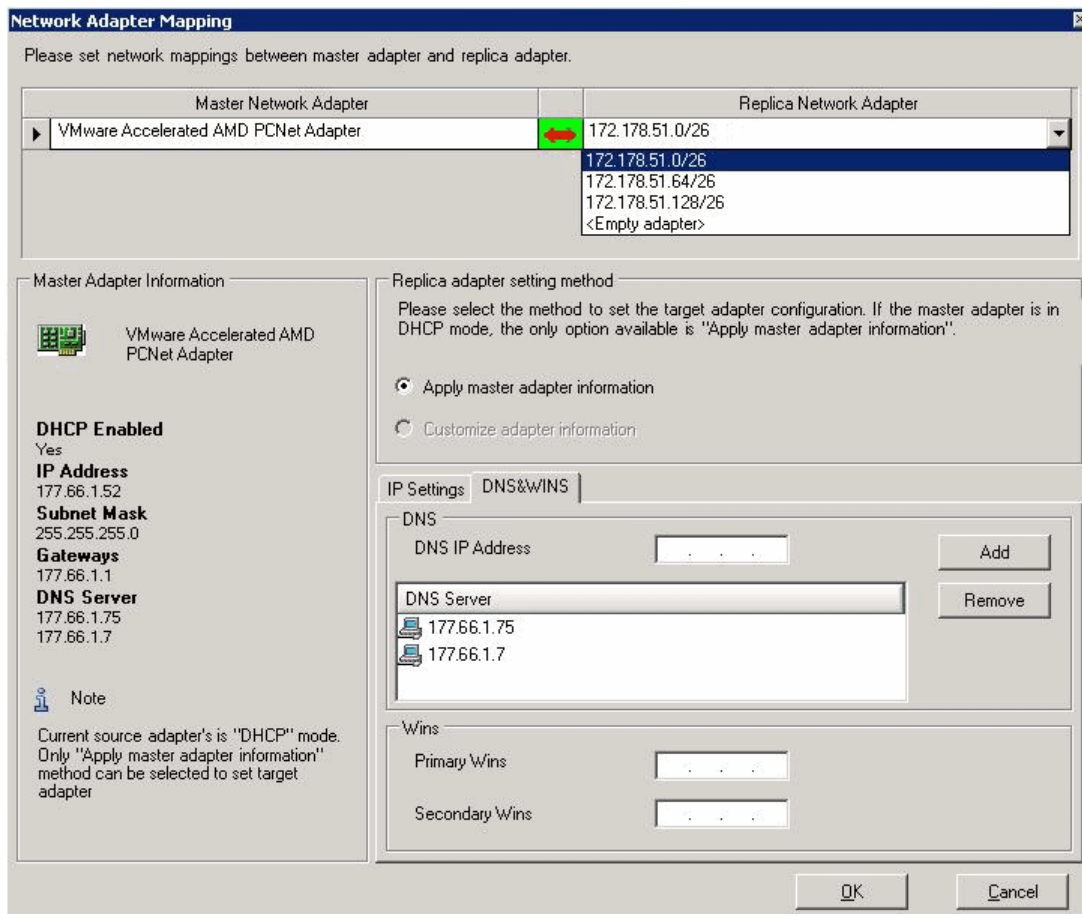
Switchover properties will be configured in this step. The recommended default values are listed below. Refer to the Administration Guide before modifying the current settings.

Property	Value
Master IPs in DNS	
IP Address	162.46.2.106
IP Address	Click here to add new IP.
User-Defined Scripts	
Physical Network Mappings	Physical network mapping assigned
Is Alive	
DB Management	
Action upon Success	

Configure mapping between the physical networks on master and replica

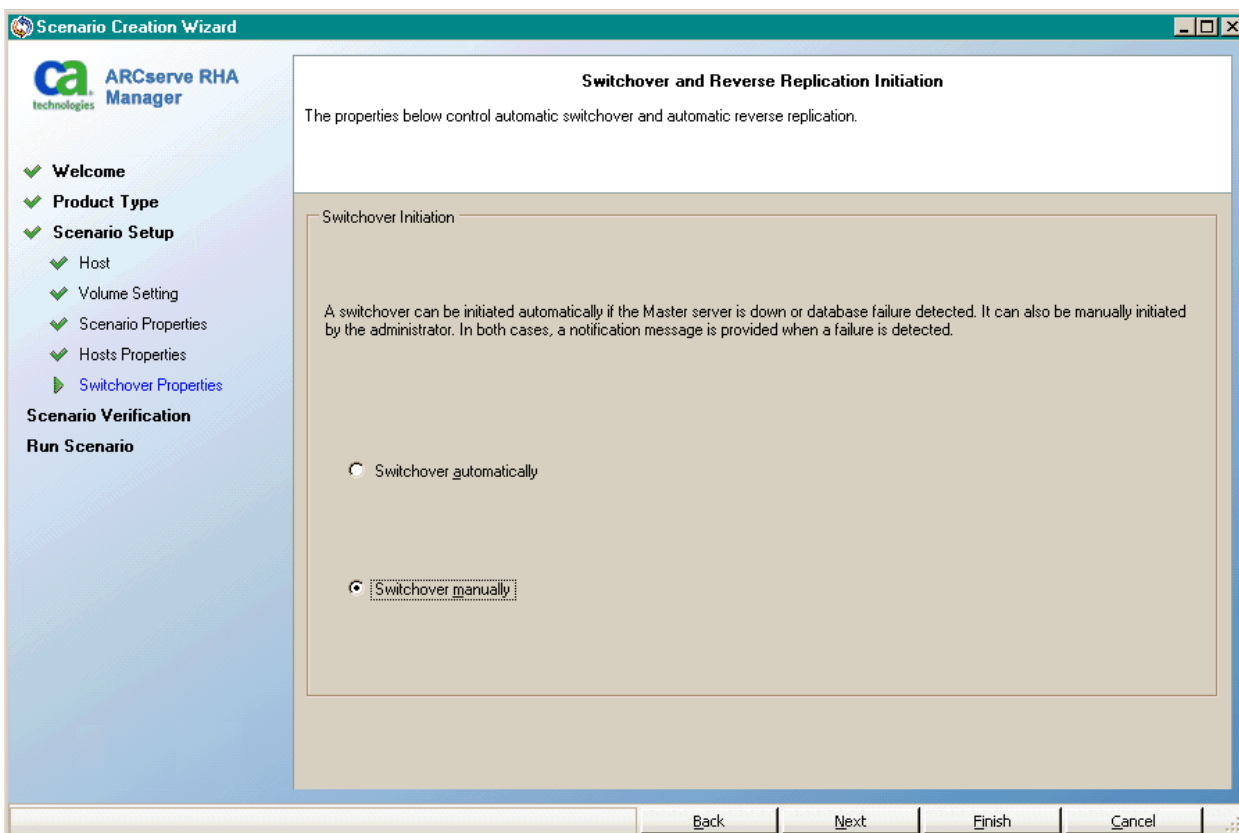
12. Expand the Physical Network Mappings option under *Network Traffic Redirection*.

The Network Adapter Mapping dialog opens.



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

The Switchover and Reverse Replication Initiation screen opens.



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

Wait for Scenario Verification to complete.

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

The Scenario Run screen opens.

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

Additional Properties for Full System Scenarios

You can manually change the following properties after the scenario creation process is complete:

Virtual Machine Local Directory (Scenario Properties tab)

In the Replica Root Directory list, double-click the virtual machine local directory to browse and select a new location. Click OK and then click the disk icon to save the modified scenario.

CPU Core (Replica Properties tab)

On the High Availability Properties tab, expand Virtual Machine settings. Click CPU Core to modify the number.

Memory Size (Replica Properties tab)

Click Memory Size to modify the value.

Virtual Network Mappings (Scenario Properties tab)

Besides the Virtual Network Mappings setting, click the entry, "Click to edit Virtual Network Mappings" to display the Physical Network Mapping dialog.

Virtual Machine Name (Replica Properties tab)

You may change the Virtual Machine Name, if desired. By default, the Master hostname is used.

Save the modified scenario when you are done modifying these settings.

Configure Additional Properties in Full System-to-ESX Scenarios

When configuring full system scenarios on ESX servers, you must manually enter an IP address on the Replica server to prevent scenarios from failing to run.

To manually enter a virtual machine IP address on the Replica server

1. Configure the scenario as described in the Create Full System Scenarios topic.
2. When you reach the Switchover Properties screen, expand the Network Traffic Redirection properties.
3. Set the following property to On: Redirect DNS.
4. Enter the DNS IP address.
5. Expand Master IPs in DNS. Click the value column and enter the IP address for the Master host.
6. Expand Virtual Machine IPs on Replica server in DNS.
7. Select Click here to add new IP and specify the IP address of the virtual machine on the Replica server to be used in case a switchover or failover is triggered.

Configure Additional Properties in Full System EC2 Scenarios

For full system EC2 scenarios, you can specify additional properties for Cloud and Virtual Machine as required for your environment.

To configure additional full system EC2 properties, configure the scenario as described in [Create a New Full System EC2 High Availability Scenario](#) (see page 123). When you reach the Master and Replica Properties screen, expand the Cloud and Virtual Machine properties and work with the following options:

Cloud Properties:

Cloud Provider

Identifies the name of the cloud service running the selected cloud instance. This is a read-only property.

Cloud Account ID

Identifies the account ID of the AWS account. This is a read-only property.

Cloud Region

Identifies the VPC region of the AWS account. This is a read-only property.

Cloud Instance ID

Identifies the ID of the cloud instance. This is a read-only property.

Cleanup Cloud Resources When Remove Scenario

Lets you specify whether to clean up cloud resources when a scenario is removed. For Full System EC2 Data Replication or High Availability scenarios, several cloud resources can be used such as the cloud instance used for failover, volumes, and snapshots. If these cloud resources are useless after a scenario is removed, you can enable this option to clean up these resources. This option is disabled by default.

Shutdown Instance on scenario stop

Lets you specify whether to shut down the replica instance automatically on scenario stop. This option is disabled by default, which means that the replica instance will not be automatically stopped if the scenario is stopped.

Virtual Machine Properties:

Virtual Platform Setting

Lets you review the settings for the following read-only properties:

Virtual Platform Type

Identifies the virtual platform type of the cloud account.

Virtual Platform

Identifies the virtual platform server of the cloud account.

Port

Identifies the port number used to connect to the virtual machine.

SSL Connection

Identifies whether the SSL (secure socket layer) connection is on or off.

Virtual Machine Setting

Lets you define the following properties:

EC2 Instance Type

Lets you specify the size assigned to the EC2 instance on the virtual machine. You can specify the appropriate instance type based on the operating system of the master and the requirements of your environment. Instance type options include:

- Small Instance
- Large Instance
- Extra Large Instance
- High-Memory Extra Large Instance
- High-Memory Double Extra Large Instance
- High-Memory Quadruple Extra Large Instance
- High-CPU Medium Instance
- High-CPU Extra Large Instance

Available options are related to the master's platform. If the master is a 32-bit operating system, only the Small Instance and High-CPU Medium Instance are available. If the master is a 64-bit operating system, then all of the other types are available.

Virtual Machine Name

Specifies the name of the virtual machine managed on the virtual platform server.

Redirection Methods for Full System Scenarios

For Full System HA scenarios, Switch computer name and IP redirection are not supported; you may optionally choose DNS redirection. You can also customize the network resource to be used on the virtual machine. Specify the IP, gateway DNS and WINS address on the mapped NIC, if it uses static IP.

During Auto-Discovery, the software obtains only the Master server IP address on the DNS server. If you set Redirect DNS to On, you must manually enter the IP address of the virtual machine on the Replica server or switchover will fail.

To manually enter the virtual machine IP Address on Replica server in DNS

1. Create the full system scenario, as usual or modify an existing scenario from the CA ARCserve RHA Manager.
2. On the Switchover Properties screen, expand the Network Traffic Redirection properties group.
3. Expand Redirect DNS properties.
4. Expand Virtual Machine IPs on Replica server in DNS.
5. For the IP Address property, click the Value column ("Click here to add new IP").
6. Specify the IP address.
7. Save modified scenario or click Next to continue creating it.

Running a Full System HA Scenario

When you start a Full System HA scenario, the software first validates the scenario configuration with a verification check. During synchronization, the data in protected volumes on the physical machine is replicated to a virtual disk file on the Replica. You can select the file, block, or volume level synchronization. Volume synchronization is recommended for the initial synchronization because it performs better in LAN-based environments. However, since volume synchronization synchronizes all data from the Master to the Replica, the file or block level synchronization is recommended for any subsequent synchronizations since they offer a significant load reduction. For resynchronization, block level sync is used by default.

During replication, the software replicates all file system changes that occur on the physical machine to the Replica and applies those changes to the data on the virtual disk.

Operations on a Virtual Machine

You can use available operations to manage a virtual machine. These operations lets you start, stop, and delete all VM resources.

Start or Stop a Virtual Machine

Use this operation to start or stop a virtual machine from its latest system status or from a bookmark. You can start or stop a virtual machine after you create a scenario and synchronize the master and replica. Use this feature when the scenario is not running. This feature is available for Full System DR and HA scenarios. The Start/Stop is a toggle menu item.

Follow these steps:

1. Click Tools, Start/Stop VM.

The Set up the virtual machine dialog opens.

2. Select a bookmark and an option to set the network for the newly created virtual machine.

Consider the following points before you use this operation:

- If you perform any of these functions: Switchover, Failover, Start VM, Recover Active Replica server, or BMR Restore, RHA automatically creates a bookmark, "bookmark for last state of the scenario". This bookmark is to keep the last data state of the scenario before the scenario stopped or performed switchover. This bookmark is internally deleted when you run the forward scenario next time.

You can go back to the last scenario state any time.

- RHA keeps the latest system status before you performed the Start VM operation. When you start the VM again, you can find a bookmark, "Latest System Status of bookmark [bookmark name]" in the bookmark list.
- When you select any other bookmark to start the VM after the VM is stopped, all data changes in latest system status is lost.

This operation creates and starts a new virtual machine on the appliance. Depending on the platform, virtual machine tools such as VMware tools, Integration Service, or XenServer tools are also installed on the virtual machine. Log in and check if data is accurate and various services are working.

Important: After the VM is started, do not use virtualization platform management tools (such as VMware vSphere Client, Hyper-V Manager, or XenCenter) to create a VM snapshot. This results in an unexpected error.

Delete Virtual Machine Resources

When you run a full system scenario, some temporary resources are created such as disk files, snapshots, and other files. This operation lets you delete these resources and is available when the scenario is not running.

Follow these steps:

1. Click Tools, Delete all VM Resources.

A warning appears that all protected data will be deleted.

2. Select Yes to clean up all VM resources.

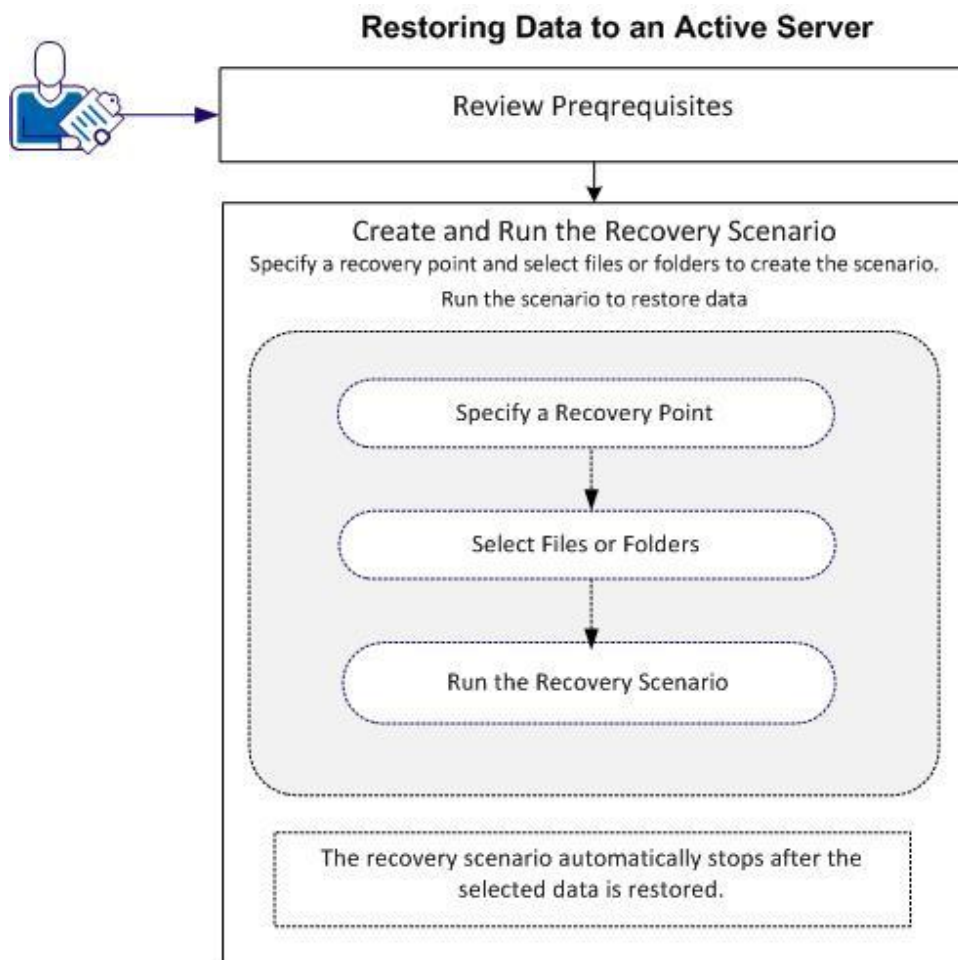
Note: If the VM is already running, the operation shuts down the VM first and then removes all VM resources.

All created image as well as disk files and snapshots are deleted.

Restore Full Systems

As an Administrator, you are responsible for ensuring that data is replicated and failover happens in case the Master server fails. On a failover or switchover, you can recover the entire data to an active server from the Replica.

The following diagram illustrates how the recovery process restores data to an active server.



Perform the following tasks to restore data:

- [Review Prerequisites](#) (see page 139)
- [Create and Run the Recovery Scenario](#) (see page 140)

Review Prerequisites

Before you can restore data to an active server, consider the following points:

- Start recovery only when the full system scenario is stopped or a switchover/failover is triggered.
- Restore data to a similar machine. Install the same version of the operating system and applications as much as possible.
- Build the new Master first and then launch the recovery process.
- For an FSP scenario, do not manually stop the virtual machine that performed a switchover or failover before you start the Master.
- Stop critical services before launching the recovery process, such as Database services and Exchange server services. When the services are running, some applications lock their files and cannot be opened during the recovery process. Start these services again after the recovery.

Create and Run the Recovery Scenario

Review the prerequisites carefully before you launch the recovery. Use the Data Recovery wizard to restore data from the Replica.

Follow these steps:

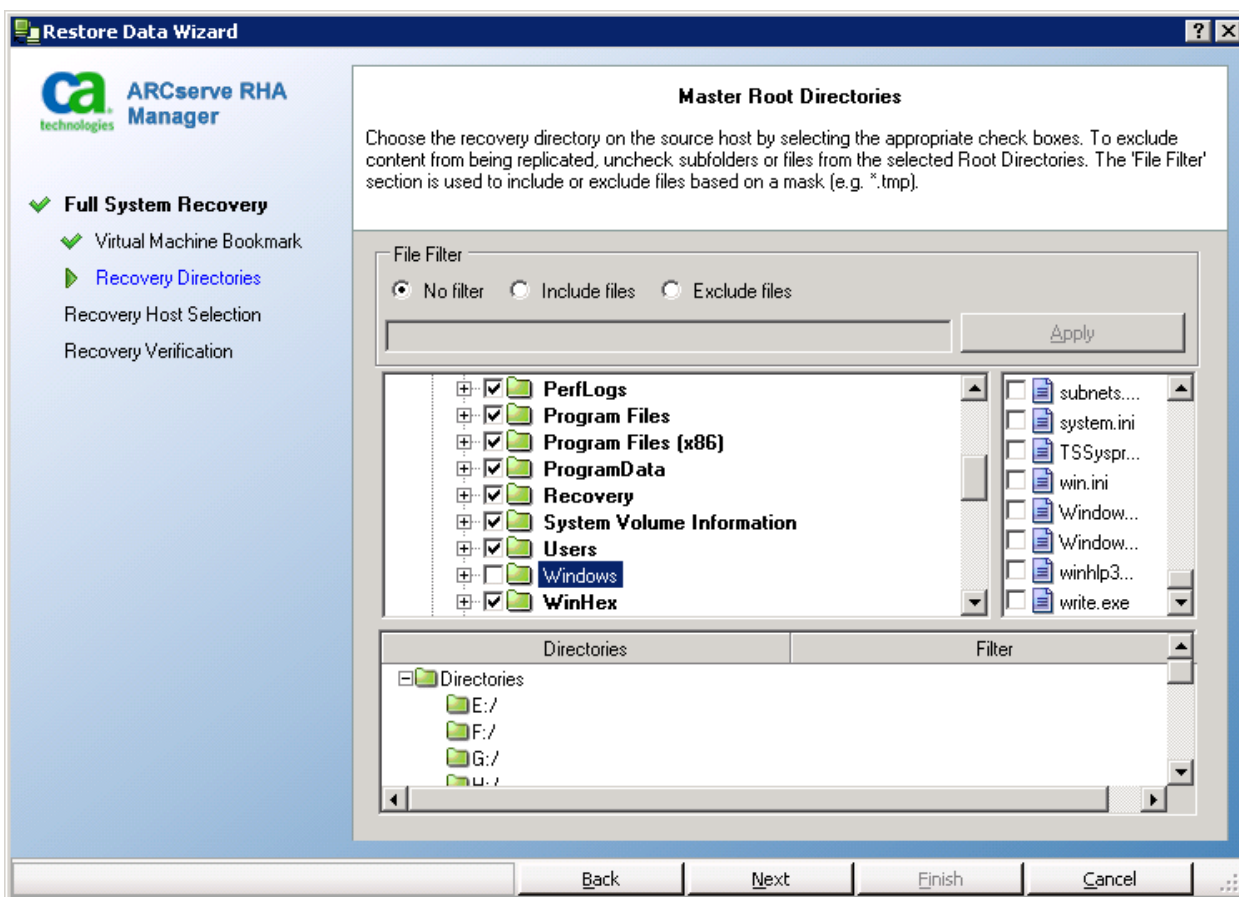
1. Log in to CA ARCserve RHA as an Administrator.
2. From the Quick Start pane, click Scenario Management to open the CA ARCserve RHA Manager.
3. Launch the Data Recovery Wizard from the Manager screen.
4. Enter the credentials, when prompted.
5. Select a recovery point on the Recovery Point Selection screen. For example, click Latest System Status and then click Next.

The virtual machine shuts down.

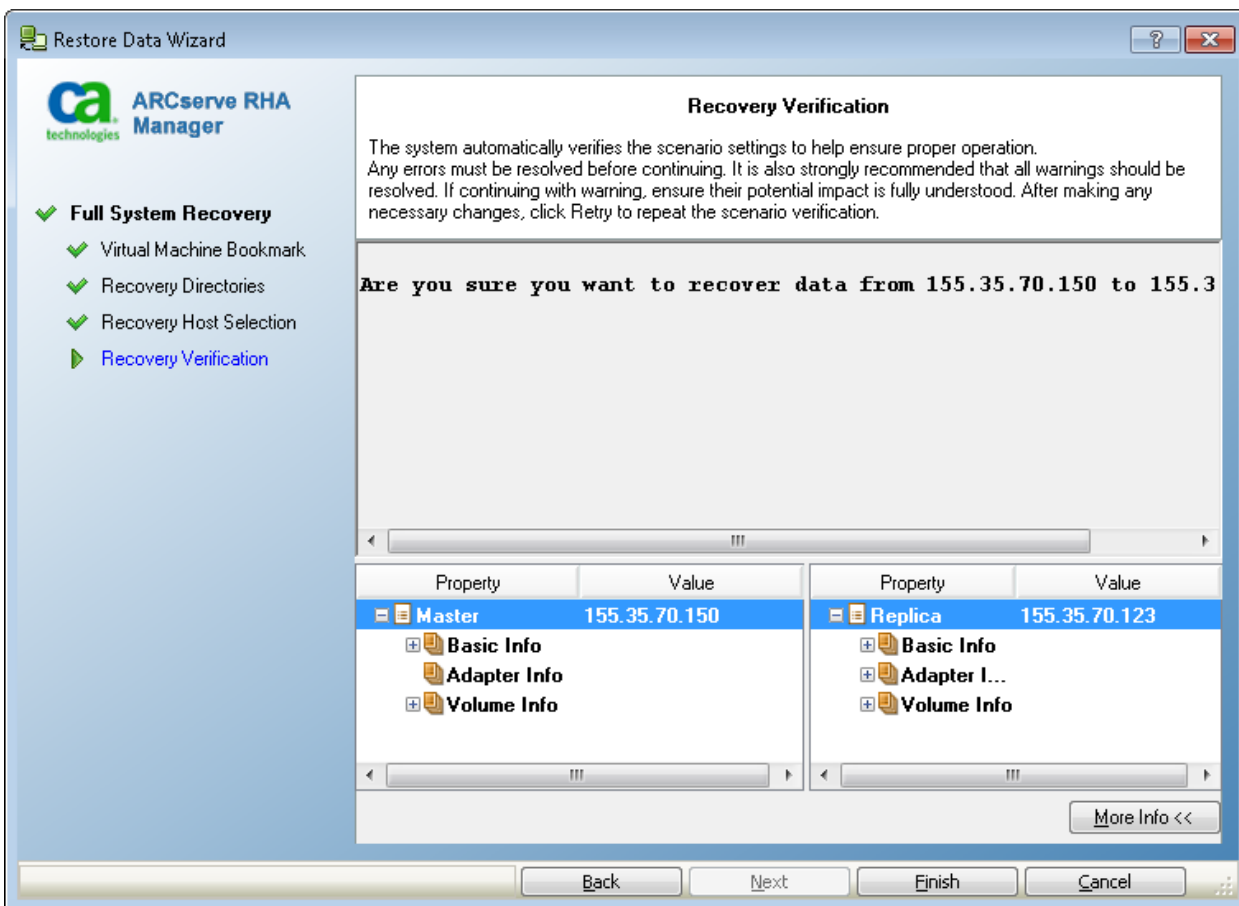
6. On the Master Root Directories screen, expand the source host and then, click or clear check boxes to include or exclude folders from the list. If you wish, you can apply a file filter. Click Next.

Important: Do not select the C:\Windows folder as it can lead to certain conflicts later, especially if the operating system versions are different. If you do select the folder, then system registry files such as Sam, Security, Software, and Default are not restored.

Note: We recommend skipping the RHA Engine folder (C:\Program Files\CA\ARCserve RHA\Engine) when you are not sure about the engine versions.



- On the Recovery Host screen, enter the required credentials and click Next.
Wait while the verification completes.



- When the Recovery Verification screen opens, review the information or click More Info to see more details. Click Finish.

A new scenario, Recovery_<Scenario Name>, is added to the list on the Manager screen. View the statistics for progress. The new scenario automatically stops once the selected data is recovered. The system reboots if necessary.

DNS Redirection using the Update DNS Tool

Use this tool to change the DNS A record of the master server to resolve to the IP address of the replica. You use this tool after you start or stop the VM using the Start/Stop VM feature. This redirection is applicable when the master and replica are on the same or different subnets. You can use this tool by running the `update_dns.exe` file in the RHA engine installation folder.

Syntax:

```
Update_dns.exe -dns -hostname -hostip -ttl -username -password -keyfile
```

Example: Update the master server address, `master.rha.com` (199.100.5.1) to that of the DNS server (199.200.2.1).

```
Update_dns.exe -dns 199.100.5.1 -hostname master.rha.com -hostip 199.200.2.1 -username test -password test
```

Example: Update the local server address to that of DNS servers 199.100.4.1 and 199.100.6.1.

```
Update_dns.exe -dns 199.100.4.1,199.100.6.1
```

Restoring Data to a Bare Metal Machine

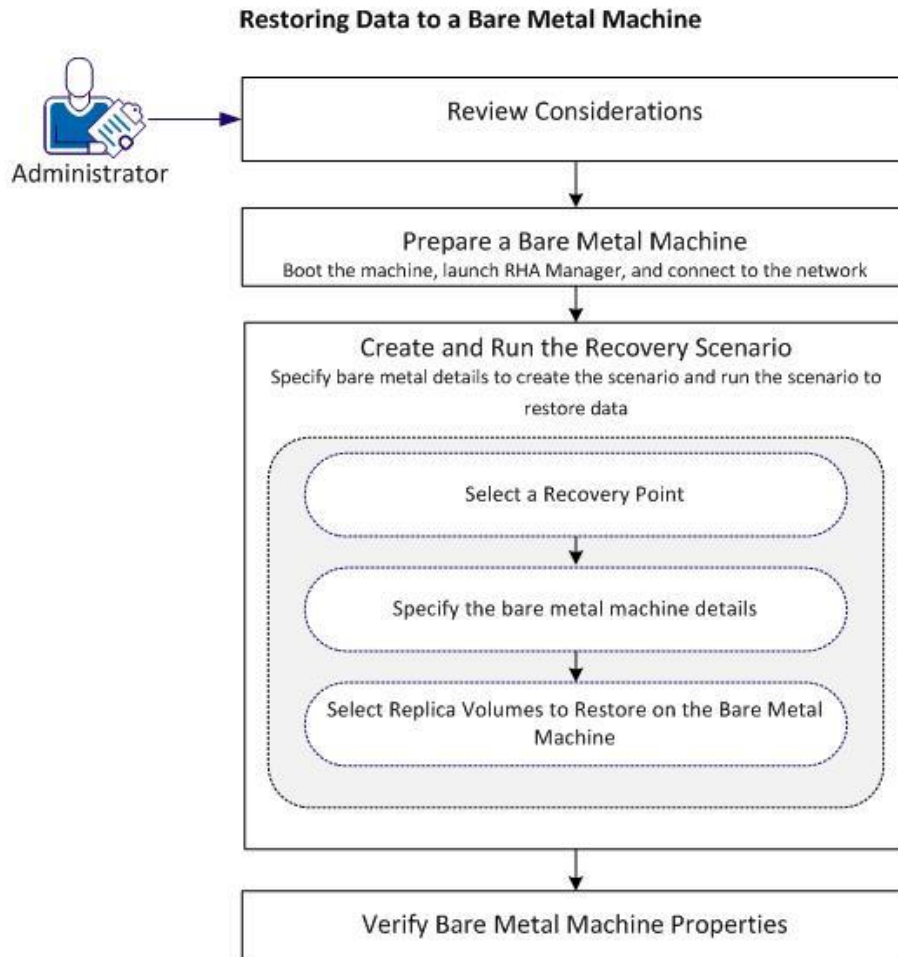
As an Administrator, you are responsible for ensuring that data is restored in case the Master server fails.

CA ARCserve RHA lets you restore data to a bare metal machine. This feature not only restores data but also builds the bare metal machine by installing the operating system, applications, and other necessary components. After you complete the process, the bare metal machine is ready to take over as the Master server.

A bare metal recovery provides the following benefits:

- Builds a machine exactly similar to the Master including operating system, applications, and data
- Eliminates manual setting up of a machine
- Recovers data to a dissimilar hardware

The following diagram illustrates how you can build and restore data to a bare metal machine from the Replica.



Perform the following tasks to restore data:

- [Review Prerequisites](#) (see page 146)
- [Prepare a Bare Metal Machine](#) (see page 146)
- [Create and Run the Recovery Scenario](#) (see page 147)
- [Verify Bare Metal Machine Properties](#) (see page 147)

Review Prerequisites

Before you restore data to a bare metal machine, consider the following points:

- Applies only to Full System Scenarios
- Supports only basic disks, no support for dynamic disks.
- Upgrade considerations: If you upgrade the RHA engine from previous releases to Service Pack 16.0.02, you must synchronize the Master and Replica at least once to complete the upgrade.

Prepare a Bare Metal Machine

Before you create a recovery scenario, prepare a bare metal machine on which you want to restore data. In this process, boot the machine with the RHA BMR CD to start the machine and launch the RHA engine. An IP address is assigned to the machine. Specify this IP address while creating the recovery scenario.

Follow these steps:

1. Insert the RHA BMR CD and boot the computer.
2. Select CA ARCserve RHA BMR to boot the system.
3. Select the language of your choice and click Next.

The RHA CD detects network adapters and disks on the bare metal machine.

4. (Optional) If the RHA CD is unable to detect network adapters and disks, click appropriate links on the left panel to install them.

The RHA engine is launched and the machine is assigned an IP address.

Create and Run the Recovery Scenario

After you prepare the bare metal machine, create a recovery scenario with the bare metal machine as the Replica. Creating this scenario lets you specify the Master and Replica (bare metal machine) details and configuration properties. When you run the scenario, data is restored to the bare metal machine.

Specify the following details on Restore Data Wizard screens to create the scenario:

- IP address of the bare metal machine.
- Volumes--Select the volumes on the Master that you want to restore. You can also manually resize the volume to a different size on the bare metal machine.

Note: You can exclude volumes that you do not want to restore.

When you run this scenario, data is restored to the bare metal machine.

Follow these steps

1. Log in to CA ARCserve RHA as an Administrator.
2. From the Quick Start pane, click Scenario Management to open the CA ARCserve RHA Manager.
3. Select the full system scenario and click Restore Data.
The Restore Data Wizard opens.
4. Follow the instructions on wizard screens to create and run the recovery scenario.

Data is restored to the bare metal machine.

Verify Bare Metal Machine Properties

After the data is restored, the bare metal machine is rebooted. Now this machine has the same name, operating system, applications, and data components as that of the original Master server. Verify these details by checking operating systems and other application properties.

Note: When the machine is rebooted, some drivers could be missing. Manually install and reconfigure these drivers to work for any hardware changes.

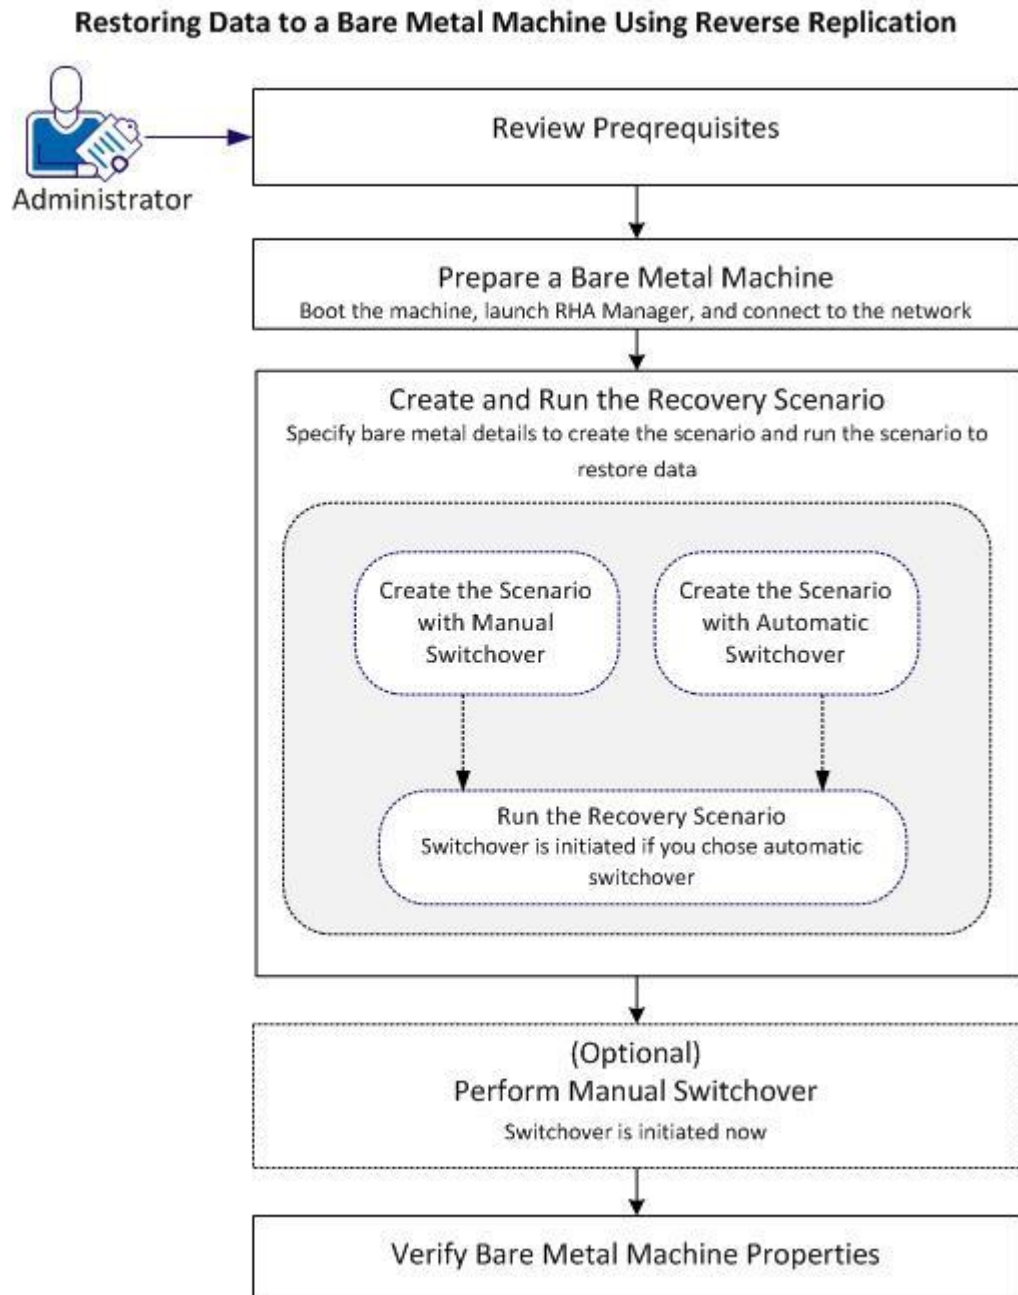
Restoring Data to a Bare Metal Machine after Failover

As a System Administrator you are responsible for ensuring that data is replicated and failover happens when a Master server does not respond. In such a case, you create a recovery scenario to restore data from its Replica. You have the option of restoring data to a bare metal machine. The bare metal recovery feature not only restores data but also builds the bare metal machine by installing the operating system, applications, and other necessary components.

A bare metal recovery provides the following benefits:

- Builds a machine exactly similar to the Master before restoring data
- Recovers data even to a dissimilar hardware
- Recovers data on a volume level

The following diagram illustrates restoring data to a bare metal machine using Reverse Replication.



Perform the following tasks to restore data:

- [Review Prerequisites](#) (see page 150)
- [Prepare a Bare Metal Machine](#) (see page 150)
- [Create and Run the Recovery Scenario](#) (see page 151)
- [\(Optional\) Perform a Manual Switchover](#) (see page 152)
- [Verify the Bare Metal Machine Properties](#) (see page 152)

Review Prerequisites

Before you restore data to a bare metal machine, consider the following points:

- Applies only to full system scenarios that have performed a failover or switchover.
- Supports basic disks and does not support dynamic disks.
- Upgrade considerations: If you upgrade the RHA engine from previous releases to Service Pack 16.0.02, you must synchronize the Master and Replica at least once to complete the upgrade.

Prepare a Bare Metal Machine

Before you create a recovery scenario, prepare a bare metal machine on which you want to restore data. Boot the machine with the RHA BMR CD to start the machine and launch the RHA engine. An IP address is assigned to the machine. Specify this IP address while creating the recovery scenario.

Follow these steps:

1. Insert the RHA BMR CD and boot the computer.
2. Select CA ARCserve RHA BMR to boot the system.
3. Select the language of your choice and click Next.

The RHA CD detects network adapters and disks on the bare metal machine.

4. (Optional) If the RHA CD is unable to detect network adapters and disks, click appropriate links on the left panel to install them.

The RHA engine starts and an IP address is assigned to the machine.

Create and Run the Recovery Scenario

Now the bare metal machine is in the network with a unique IP address. To restore data, create a recovery scenario with this bare metal machine as the Replica. This way, you can restore data from the virtual machine to the bare metal machine.

Specify the following details on Restore Data Wizard screens to create the scenario:

- IP address of the bare metal machine.
- Volumes--Select the volumes on the Master that you want to restore. You can also manually resize the volume to a different size on the bare metal machine.
Note: You can exclude volumes that you do not want to restore.
- Network Adapters—Select the Replica network adapter that you want to map to the Master network adapter. For servers in a workgroup, use Redirect DNS. To use Redirect DNS, specify domain administrator credentials.
- Switchover type--Automatic or Manual.

When you run this scenario, data is restored to the bare metal machine. Depending on the switchover type, either the switchover process is automatically triggered or you manually initiate the process.

Follow these steps

1. Log in to CA ARCserve RHA as an Administrator.
2. From the Quick Start pane, click Scenario Management to open the CA ARCserve RHA Manager.
3. Select the full system scenario that performed switchover or failover and click Run.

The Restore Data Wizard opens.

4. Follow the instructions on wizard screens to create and run the recovery scenario.

Data is restored to the bare metal machine. If you selected automatic switchover, the switchover process is initiated and the bare metal machine is ready to take over as the Master server.

Note: If you selected the manual switchover, manually initiate the switchover process. See (Optional) Perform a Manual Switchover.

The switchover process performs the following tasks:

- Moves the resources like IPs, DNS from the Master server to the bare metal machine.
- Performs driver injection on the bare metal machine.
- Shuts down the Master.
- Reboots the bare metal machine.

Note: if you stop a scenario, then above tasks are not executed and bare metal machine is not recovered completely.

(Optional) Perform a Manual Switchover

If you chose manual switchover, data is synchronized after you run the recovery scenario. To perform the switchover, select the scenario and manually initiate the switchover process.

Follow these steps

1. Open CA ARCserve RHA Manager and select the recovery scenario from the Scenario pane. Verify that the scenario is running.

2. Click Perform Switchover.

A confirmation message appears.

3. Click Yes.

A switchover process is initiated and the Master server is disconnected. The bare metal machine is rebooted.

Now, the bare metal machine is ready to take over as the Master server.

Verify Bare Metal Machine Properties

After the switchover process is complete, verify that the machine name, operating system, applications, and data components are same as that of the original Master server.

Note: When the machine is rebooted, some drivers may be missing. Manually install these drivers.

Now if you want to run the original forward scenario, perform the Recovering Active Server process to set the recovered Master as the Active Server. See Recover Active Server.

Perform Full System Assured Recovery Testing

You can perform Assured Recovery Testing for Full System HA scenarios manually or automatically.

During AR Testing, the software suspends journal application, creates a VM with the virtual hard disk file and a snapshot to save possible changes, then starts the VM. If you chose automatic AR testing, VMware tools are installed. Wait for the software to stop the test. If you chose manual, you'll need to stop the AR test manually.

Once stopped, the software resumes applying journal changes.

For more information, see the section, [Testing Assured Recovery and Managing VSS Snapshots](#), in the CA ARCserve RHA Administration Guide.

Recover Active Server for Full System Scenarios

The Recover Active Server process lets you forcibly make the Master or Replica server the active server without completing the data synchronization process. This action is necessary in cases where you do not want to synchronize data. For example, if switchover occurred but no data was changed on the Replica server. In this case, you can even have newer data on the Master server making it undesirable to synchronize data from the Replica to the Master server.

To use this option, make sure that the scenario is stopped, and select Recover Active Server from the Tools menu. Select either Make Master Active or Make Replica Active depending on which server you want to make active. If you make the Master active, the process stops and unregisters the virtual machine (that the appliance created). If you make the Replica active, the appliance registers and starts the virtual machine.

Note: When the platform is Amazon EC2, the appliance would stop the VM without unregistering. Register or unregister the VM separately.

Important! While this option is the right choice in many situations, use it with caution. When used improperly data loss can occur. Usually, RHA does not allow switching from one host to another until all data is synchronized. When using Recover Active Server, 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, manually make sure 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 34).

The following scenarios are some examples where you can use the recover active server process.

Scenario 1: Making the Master server active after a switchover or failover

After switchover or failover, the Master server becomes the standby server, and the Replica (virtual machine the appliance created) the active server. Now if you do not want to make the Replica as the active server, then you can use this feature and can make the master active.

Follow these steps:

1. Click Tools, Recover Active Server.
The Recover Active Server dialog opens.
2. Click Make Master Active.
3. The appliance unregisters the virtual machine.

You can run the forward scenario again.

Scenario 2: Making the Replica server active after synchronization

You created a scenario and just synchronized the Master and Replica. In this case, the master is active and the replica is inactive. In some cases, you want to make the replica active, for example, if you want to check whether the replica (virtual machine) is available and working as expected.

Follow these steps:

1. Click Tools, Recover Active Server.
The Recover Active Server dialog opens.
2. Click Make Replica Active.
The appliance registers and starts the virtual machine.

Additional Tools

This section provides you with details on the following tools:

- Virtualization hypervisor login setting tool
- Update DNS tool

Virtualization Hypervisor Login Setting Tool

Use this tool to set up the login credential information for a specific virtualization hypervisor on an existing full system scenario. Each login credential record is uniquely identified by its hostname/IP and the port. To use this tool, run the `ws_p2v_log_in_setting.exe` file from the RHA engine installation folder on the appliance server and follow the instructions.

Index

(

(Optional) Perform a Manual Switchover • 150, 152

A

About This Guide • 9
Add a New Cloud Account • 109
Add CA-IP to Existing Scenarios • 47
Add CA-IP to New Scenarios • 48
Add IP on the Master Server • 45
Additional Information and Tips • 89
Additional Properties for Full System Scenarios • 131
Additional Tools • 154
Automatic Redirection Using Switch Computer Name • 49

C

Change the Default AWS Region • 111
Configure Additional Properties in Full System EC2 Scenarios • 133
Configure Additional Properties in Full System-to-ESX Scenarios • 104, 132
Configure HTTP Communication Protocol on vCenter Server 4.0 Systems • 99
Configure Hyper-V for High Availability • 64
Configure the vCenter Server Database Server on the Replica • 17, 18
Configure the vCenter Server License Server on the Replica • 17, 20
Configure the vCenter Server Master Server • 16
Configure the vCenter Server Replica Server • 16
Configure the Web Proxy to Connect to the Cloud Service • 108, 110
Configuring Master and Replica for Full System Scenarios • 98
Considerations for Hyper-V scenarios across WAN • 69
Create a New EC2 Data Replication Scenario • 119
Create a New EC2 Replica Instance • 112
Create a New Full System EC2 High Availability Scenario • 123, 133
Create a New Hyper-V High Availability Scenario • 75
Create a New Hyper-V Replication Scenario • 70
Create and Run the Recovery Scenario • 138, 140, 145, 147, 150, 151

Create Full System Scenarios • 99
Create Full System Scenarios for Hyper-V Platform • 100, 105
Create vCenter Server HA Scenarios • 23
Create vCenter Server Replication Scenarios • 21

D

Data Rewind • 31
Delete a Cloud Account • 111
Delete an EC2 Replica Instance • 118
Delete Virtual Machine Resources • 137
DNS Redirection • 45
DNS Redirection using the Update DNS Tool • 143
Documentation Changes • 4

H

How Full System Scenarios Work • 95
How Redirection Works • 44
How Switchover and Switchback Work • 51, 77
How to Restore Data on Hyper-V Machines • 87
How to Restore Data on vCenter Machines • 33
How to Start a Hyper-V VM on the Replica Server • 89
Hyper-V HA Configuration • 62
Hyper-V HA Properties • 72
Hyper-V Replication and High Availability • 70
Hyper-V Replication Properties • 71
Hyper-V Server Configuration Requirements • 61
Hyper-V VM Auto-Discovery • 63

I

Initiate Switchback • 54, 80
Initiate Switchover • 51, 52, 77, 78
Introduction • 9

L

Log On Account Conditions • 10, 69

M

Manage Services • 92
Manually Recover a Failed Server-Move IP Address • 34
Manually Recover a Failed Server-Switch Computer Name • 35

Manually Recover Failed Server-IP and Switch
Computer Name • 34, 36
Move IP Redirection • 45

O

Operations on a Virtual Machine • 136

P

Perform Full System Assured Recovery Testing • 105,
153
Prepare a Bare Metal Machine • 145, 146, 150
Product References • 3
Protecting Full Systems • 95
Protecting Hyper-V Environments • 61
Protecting VMware vCenter Server Environments •
13

R

Recover Active Server • 29
Recover Active Server for Full System Scenarios •
153
Recover Hyper-V Data with Rewind Points • 88
Recover Lost Data from Replica • 28
Recovering Servers • 29, 34, 153
Redirection Methods • 44
Redirection Methods for Full System Scenarios • 135
Register CA ARCserve RHA Licenses • 10
Related Documentation • 10
Restore Full Systems • 138
Restoring Data to a Bare Metal Machine • 144
Restoring Data to a Bare Metal Machine after
Failover • 148
Review Prerequisites • 138, 139, 145, 146, 150
Run a Scenario • 56
Run a Scenario from Outside the Wizard • 83
Running a Full System HA Scenario • 135

S

Scripts Redirection • 49
Set Scenario Properties • 23, 24
Setting Bookmarks • 30, 86
Spool Directory Settings • 26, 72, 74, 91
Start an EC2 Replica Instance • 118
Start or Stop a Virtual Machine • 136
Stop a Scenario • 58, 84
Stop an EC2 Replica Instance • 118
Supported Databases VMware vCenter Server • 19
Switch Computer Name Redirection • 48

Switching Over and Switching Back • 24, 50, 76
Switchover Considerations • 56, 82

T

The Data Recovery Process • 27, 86
Troubleshooting Hyper-V • 89
Troubleshooting Oracle Databases • 44
Troubleshooting vCenter Server Scenarios • 36

U

Update Cloud Account Information • 111

V

vCenter Server Replication and High Availability • 21
Verify Bare Metal Machine Properties • 145, 147,
150, 152
View a Report • 58, 84
Virtualization Considerations • 11
Virtualization Hypervisor Login Setting Tool • 155
VMware vCenter Server RHA Configuration • 14