

Microsoft Exchange Server Operation Guide

Arcserve® Continuous Availability

Version 18.0

arcserve®

Legal Notices

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

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

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

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

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

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

The manufacturer of this Documentation is Arcserve.

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

© 2020 Arcserve, including its affiliates and subsidiaries. All rights reserved. Any third party trademarks or copyrights are the property of their respective owners.

Arcserve Product References

This document references the following Arcserve products:

- Arcserve® Continuous Availability
- Arcserve® Live Migration
- Arcserve® Assured Recovery®
- Arcserve® Content Distribution

Contact Arcserve

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

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

With Arcserve Support:

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

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

Providing Feedback About Product Documentation

If you have comments or questions about Arcserve product documentation, please contact [us](#).

Contents

Chapter 1: Introduction	7
Support for New Microsoft Exchange Server Features	8
About This Guide	9
Related Documentation	10
Chapter 2: Protecting Microsoft Exchange Server Environments	11
Base Configuration	12
Exchange Server Configuration Requirements	13
Configuring Customized Service Management	14
Modify the Exchange Configuration on the Master and Replica	15
Log On Account Conditions	16
Create an Exchange Server Domain User Account	17
About Clusters	19
Register Arcserve Continuous Availability Licenses	20
Chapter 3: Creating Replication and High Availability Scenarios	23
Create an Exchange Replication and Data Recovery Scenario	24
Create an Exchange High Availability Scenario	26
Chapter 4: Managing Replication and High Availability Scenarios	29
Set Scenario Properties	30
Manage Services	33
Run the Scenario from Outside the Wizard	36
Stop a Scenario	38
View a Report	39
Chapter 5: Switching Over and Switching Back	41
How Switchover and Switchback Work	42
Initiate Switchover	44
Initiate Switchback	45
Switchover Considerations	47
Chapter 6: Recovering Data	49
The Data Recovery Process	50
Setting Bookmarks	51
Data Rewind	52
Recover Lost Data from Replica	56
Chapter 7: Redirection Methods	61
How Redirection Works	62

DNS Redirection	63
Move IP Redirection	64
Add IP on the Master Server	65
Add Arcserve-IP to Existing Scenarios	67
Add Arcserve-IP to New Scenarios	68
Cluster Move IP	69
Use the Master Cluster	70
Use the Manager	71
For New Scenarios	72
For Existing Scenarios	73
Scripts Redirection	74
Mailbox Redirection	75
Disable Redirection	76
Chapter 8: Additional Information and Tips	77
Spool Directory Settings	77
Recover Active Server	77
Recovering Servers	79
Manually Recover a Failed Server-Move IP Address	79

Chapter 1: Introduction

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

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

- You are familiar with Arcserve Continuous Availability 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:

Support for New Microsoft Exchange Server Features	8
About This Guide	9
Related Documentation	10

Support for New Microsoft Exchange Server Features

Arcserve Continuous Availability supports the following for Microsoft Exchange Server:

- Auto-discovery
- Replication
- High Availability
- Data Rewind
- Assured Recovery

Microsoft introduces changes to the database hierarchy from Exchange 2010. The concept of a storage group was removed, and public folder databases and database management were moved to the organization level. With Exchange Server 2010 and 2013, database protection is provided by Database Availability Groups (DAG). Database Availability Groups are not supported in Replication or HA scenarios. When a Master or Replica is part of a DAG, the software displays a warning.

Note: CDP Repository is no longer supported.

Important! Due to this hierarchy change, you cannot create a database with the same name on the Master and Replica servers, even if the database is dismounted. To overcome this limitation, functionality has been built into the software that allows it to temporarily rename the database for switchover purposes.

With Arcserve Continuous Availability, replication and high availability protection was moved from the storage group level to the mailbox store level.

You can now connect to their mailbox through a server that fills the Client Access Server (CAS) role, while Exchange Server 2007 allowed users to connect directly to their mailbox server. In Arcserve Continuous Availability scenarios, the CAS server must be available in order to fulfill a client access request. The method you used to deploy CAS determines how the software protects the environment.

When the Master server also fulfills the CAS role, the Replica must do the same. You can decide at scenario creation time whether to confirm CAS on the Replica. If the Master server holds only the Mailbox role, CAS must be confirmed on the Replica at scenario creation time. When the master and replica are Exchange 2013 with CAS, you should configure NLB for the master and replica.

About This Guide

This document describes how you can implement a Arcserve Replication and High Availability solution for Microsoft Exchange Server. It is essential that you have the appropriate resources and permissions to carry out each task.

Related Documentation

Use this Guide with the *Arcserve Continuous Availability Installation Guide* and the *Arcserve Continuous Availability Administration Guide*. For links of Arcserve Continuous Availability 18.0 documentation, refer to the following:

- [Bookshelf](#)
- [Release Notes](#)

Chapter 2: Protecting Microsoft Exchange Server Environments

This section contains the following topics:

Base Configuration	12
Exchange Server Configuration Requirements	13
Configuring Customized Service Management	14
Modify the Exchange Configuration on the Master and Replica	15
Log On Account Conditions	16
Create an Exchange Server Domain User Account	17
About Clusters	19
Register Arcserve Continuous Availability Licenses	20

Base Configuration

Base Configuration

- Two servers running supported Windows Server with the same level of service packs and hot fixes installed.

Note: For a complete list of supported operating systems and applications, see the Arcserve Continuous Availability Release Notes.

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

Exchange Server Configuration Requirements

Important! For Exchange scenario to work correctly on Windows 2008 or 2008 R2, Microsoft .Net framework 4 or above must be installed on both master and replica server. For details, refer to the [KB article](#).

- Microsoft Exchange Server installed on each server. Both should have the same Exchange edition and version.
- Both servers should have identical service packs and hot fixes.
- [For Exchange Server 2010/2013] Both servers should have the Mailbox role installed. Only the Mailbox role is protected in Arcserve Continuous Availability scenarios. If the Master is not also filling the CAS role, the Replica server must fill this role. If the Master server is the only server in the Entire Exchange Organization that fills the CAS and HUB transfer roles, then the Replica server should have the identical roles installed.
- [For Exchange Server 2007/2010/2013] Both servers should have Mailbox role installed. If the Master and Replica servers are located on different sites, and there is only one Exchange server on the Replica site, both servers (Master and Replica) should have identical Exchange Server roles.
- [For Exchange Server 2007/2010/2013] Both servers should have identical PowerShell version.
- Both servers should have the same Exchange Administrative Group.
- [For Exchange Server 2013] When the Exchange Server has CAS, then you should configure Network Load Balancing. After configuring the Network Load Balancing, enter the following command:

`Get-OutlookAnywhere -server <server> | set-outlookanywhere -internalhostname <fqdn> -internalclientsRequireSSL $false`

Services Managed in Exchange Server 2007/2010

MSExchangeIS

Microsoft Exchange Information Store

MSExchangeSearch

Microsoft Exchange Search Indexer

Example:

```
Net STOP <service name>
```

```
Net STOP MSExchangeIS (stops the Microsoft Exchange Information Store service)
```

Configuring Customized Service Management

The CAS role must be verified before running scenarios. If it is not available, clients cannot access mailbox roles. You can configure your scenarios to manage the CAS role using customized service management.

To access the Customize Service Management dialog, expand the DB Management properties group on the Switchover Properties screen during scenario creation. You can also manually change DB Management properties from the High Availability tab of the Manager.

Important! The services you specify are considered critical. If any service you specify in Customize Service Management is not running, switchover is triggered.

Modify the Exchange Configuration on the Master and Replica

When the HA scenario is running, do not make any changes to the Exchange configuration on the Master (such as adding a new database to a replicated Exchange Server 2007 storage group). You must stop the scenario before making changes.

Important! If you make changes to an HA scenario that is running, after a switchover, the Exchange Server on the Replica may not be able to start.

To modify the Exchange configuration on the Master or Replica

1. Stop the scenario.
2. Make the changes on the Exchange Server on the Master.
3. Make the same changes on the Exchange Server on the Replica.
4. Run auto-discovery. (On the Framework pane, open the Root Directories tab for the Master, and double-click the Auto-discovered icon.) Auto-discovery identifies the changes made to the Master configuration automatically.
5. Save the scenario by clicking the Save button.
6. Run auto-configuration on the Replica so that it can remain consistent with the Master.
7. Restart the scenario.

Alternatively, you can stop and auto-configure an existing scenario after auto-discovery completes.

Note: Auto-discovery and auto-configuration does not synchronize Exchange or Database properties. You must apply the same changes on the Replica server.

Log On Account Conditions

The Arcserve Continuous Availability 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.

Create an Exchange Server Domain User Account

To avoid Exchange Server auto-configuration problems, make sure you are using a domain admin account. Operations fail if the local system account is used. If you cannot permit use of the Domain Admin account, use this procedure.

Note: The Network Traffic Redirection method you choose also requires permission in order to complete the redirection process. Exchange Server scenarios typically use DNS or Move IP Address redirection methods.

1. Create a Domain User account. This account will be used as a service account for Arcserve Continuous Availability. Set the password to Never Expire. If your policy is to periodically change passwords, do so manually to avoid breaking scenarios when passwords expire.
2. Assign the Arcserve Continuous Availability Engine service account to the Local Administrators Group on both the Master and Replica servers. If you do not grant the Engine service account Local Administrator privileges, you must grant the Engine service account Full Access to each directory containing data to be replicated on both the Master and Replica servers.
3. Assign the newly created service account to the Arcserve Continuous Availability Engine service on the Master and Replica servers.
 - a. Click Start, Settings, Control Panel, Administrative Tools, LocalSecurityPolicy.
 - b. Open Local Policies.
 - c. Select User Rights Management.
 - d. Find Log on as a Service.
 - e. Right-click Log on as a Service and go to Properties.
 - f. Confirm the Engine service account is listed. To add it, click Add User or Group.
 - g. In the Select Users or Groups field, make the From This Location is set to the Domain and add the Engine service account.
 - h. Click OK to close the Add User or Group dialog.
 - i. Click OK to close the Log On as a Service Property dialog.
 - j. Repeat this procedure on all servers involved in the scenario.

4. Grant the Engine Service Exchange Full Administrator privileges.
 - a. Open Exchange System Manager and select the Exchange Domain.
 - b. Choose Action, Delegate Control.
 - c. In the Exchange Administration Delegation wizard, click Next.
 - d. Click Add.
 - e. Click Browse.
 - f. Change the location to be the Domain.
 - g. Enter the name of the Engine service account.
 - h. Click OK to add the account.
 - i. Click OK in the Delegate Control Box.
 - j. Click Next to finish the Exchange Administration Delegation wizard.
5. Assign the Engine service account the appropriate permissions to the Engine service account User Object.
 - a. Open ADSI Edit.
 - b. Connect to the domain.
 - c. Open the OU containing the User Objects. By default, this is CN=Users.
 - d. Find the Engine service account object. CN=Arcserve Continuous Availability Engine service account
 - e. Right-click the object and select Properties.
 - f. Click Security tab.
 - g. Click Add.
 - h. Add the Engine service account.
 - i. Set Permissions to Full Control.
6. Assign full permission to the A or Host record of the Master server record in DNS zone.
 - a. Right-click the Master's A record and click Properties.
 - b. Click Security.
 - c. Choose Full Control rights for the Arcserve Continuous Availability service account.

About Clusters

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

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

Note: On Exchange Server 2007, LCR and SCC deployments are supported, but CCR and SCR deployments are not supported. Exchange Server 2010 no longer supports CCR, LCR, SCC and SCR.

Register Arcserve Continuous Availability Licenses

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

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

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

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

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

To register Arcserve Continuous Availability using the license key

1. Open the Manager.

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

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

The Register Arcserve Continuous Availability dialog opens.

4. Complete the following fields:

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

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

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

Chapter 3: Creating Replication and High Availability Scenarios

This section describes how to create and configure replication and high availability scenarios for Microsoft Exchange Server.

This section contains the following topics:

Create an Exchange Replication and Data Recovery Scenario	24
Create an Exchange High Availability Scenario	26

Create an Exchange Replication and Data Recovery Scenario

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

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

To create a exchange replication and data recovery scenario

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

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

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

The Master and Replica Properties dialog opens.

8. Accept default settings or make the desired changes and click Next.
9. If you selected Integrity Testing for Assured Recovery, the dialog opens now. Set a schedule if desired. For more information, see the *Arcserve Continuous Availability Administration Guide*.
10. Click Next to initiate scenario verification. If errors are reported, you should resolve them before continuing. If either the Master or Replica servers are part of a DAG, you will receive a warning message. The software also verifies that the active and standby server configurations are compatible, required services are running, and Active Directory configuration is correct. You should also ensure that the CAS server role is available. At successful verification, click Next to complete scenario creation.
11. Choose Run Now or Finish, as desired. Run Now starts synchronization. Finish allows you to run the scenario later. See [Run the Scenario from Outside the Wizard](#).

Create an Exchange High Availability Scenario

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

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

To create an Exchange Server high availability scenario

1. From the Arcserve Continuous Availability Manager, click Scenario, New, or click the New Scenario button.
2. When the Welcome dialog opens, select Create New Scenario and click Next.
3. When the Select Scenario Type dialog opens, select Exchange, High Availability (HA) Scenario, and Replica Integrity Testing for Assured Recovery (optional). For more information about Assured Recovery, see the *Arcserve Continuous Availability Administration Guide*.
4. When the Master and Replica hosts dialog opens, name your scenario and provide the hostname or IP address for the Master and Replica servers. If either server is an MSCS cluster, enter the cluster resource virtual server name or IP address. Click Next. For more information, see [Redirection Methods](#).
5. Wait for Engine Verification to complete. If needed, click Install to upgrade the Engine on one or both servers and then click Next.
6. The Database for Replication dialog opens, listing all auto-discovered results for the specified Master. By default, all mailbox stores are included. Change selections if desired and then click Next.

Important! If the replicated Exchange Server 2007 scenario has a copy of a public database and a role of Offline Address Book or Schedule+ Free Busy master required by Outlook 2003 and older clients, then include the public database in the replication scenario.

7. From the Replica Configuration dialog, click Next to automatically configure Exchange on the Replica to match the Master.

Arcserve Continuous Availability auto-configuration component verifies that the Exchange Server configuration on the Master and Replica servers are identical during the replication procedure. If there are discrepancies, Arcserve Continuous Availability performs the required actions: deletes storage groups, public folders or mailbox stores from the Replica, creates new ones and modifies existing ones. The actions performed during the configuration process are indicated in the Action column on the right.

With Exchange Server 2010/2013, mailbox database names cannot be identical. The Replica database name is shown as <MasterName>_<XXXX>, where <XXXX> represents a random number. For Exchange Server 2013, the database name does not change during switchover.

You can perform these auto-configuration tasks:

- **Create** - a new, storage group, public folder or mailbox store is created.
 - **Keep** - the existing storage items remain the same.
 - **Remove** - the existing storage items store is deleted.
 - **Update** - the existing storage items remain, but its location is changed.
8. Click Next when done.
 9. When the Scenario Properties dialog opens, configure additional properties, if needed. Click Next. For more information, see Scenario Properties or the *Arcserve Continuous Availability Administration Guide*.
 10. If you are running Exchange Server 2010, set the following property in the ws_rep.cfg file: ExDisableRedirectDNS=True (default setting). For more information, see the topic, [Disable Redirection](#).
 11. Enable the DNS redirection in the following cases:
 - ◆ The server is Microsoft Exchange 2013
 - ◆ The CAS role is used for the master
 - ◆ NLB is not correctly configured
 12. The Master and Replica Properties dialog opens. Accept default settings or make the desired changes and click Next.
 13. Wait for the Switchover Properties dialog to retrieve information. Configure the desired redirection properties and click Next. For more information, see [Switching Over and Switching Back](#).
 14. From the Switchover and Reverse Replication Initiation dialog, choose automatic or manual switchover, and automatic or manual reverse replication, as needed.

We recommend that you do not set both options to Automatic. For more information, see Scenario Properties or the *Arcserve Continuous Availability Administration Guide*.

15. Click Next to initiate scenario verification. If errors are reported, resolve them before continuing. At successful verification, click Next to complete scenario creation.
16. Choose Run Now or Finish, as desired. Run Now starts synchronization. Finish allows you to run the scenario later. See [Run the Scenario from Outside the Wizard](#).

Chapter 4: Managing Replication and High Availability Scenarios

This section contains the following topics:

Set Scenario Properties	30
Manage Services	33
Run the Scenario from Outside the Wizard	36
Stop a Scenario	38
View a Report	39

Set Scenario Properties

You can modify the scenario, using the Properties pane.

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

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

The screenshot displays the Arcserve Continuous Availability Manager interface. On the left, the 'Scenarios' pane shows a tree view with 'Exchange' expanded. Under 'Exchange', several scenarios are listed, including 'Exchange-DR', 'Exchange-HA', and 'Exchange-MoveIP'. The 'Exchange-MoveIP' scenario is selected, showing its state as 'Stopped by... HA/AR'. Below the scenario list, a table shows hosts and their status. On the right, the 'Properties' pane is open, showing the 'Scenario Properties' tab. The 'General' tab is selected, displaying a table with 'Property' and 'Value' columns. Below the properties pane, the 'Events' log is visible, showing a list of events with 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-MoveIP	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.

Settings on the Properties Tab

Scenario Properties

These settings establish default behavior for the entire scenario.

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

Master and Replica Properties

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

- ◆ Host connection properties -- Enter the IP address, Port number and Fully Qualified Name of the Master and Replica.
- ◆ Replication properties -- These properties differ for Master and Replica. See the *Arcserve Continuous Availability Administration Guide* for more information.
- ◆ Spool properties -- Set the size, minimum disk free size and directory path. See [Spool Directory Settings](#) for more information.
- ◆ Event notification properties -- specify a script to run, select email notification, or write to event log.
- ◆ Report properties -- select synchronization or replication reports, specify distribution or script execution.
- ◆ (Replica) Scheduled Tasks -- set or suspend tasks, including Replica Integrity Testing for Assured Recovery. For more details, see the *Arcserve Continuous Availability 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, or User-defined scripts.
- ◆ Is Alive properties -- set the heartbeat frequency and check method.
- ◆ DB Management properties (does not apply to File Server scenarios) -- instructs Arcserve Continuous Availability to manage shares or services on a database server.
- ◆ Action upon Success properties -- defines custom scripts and arguments for use.


Manage Services

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

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

The following steps are for Custom Application scenarios.

To manage services



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

- **Microsoft SharePoint Server** - lists SharePoint Server-related services if the current host has SharePoint Server installed
- **VMware vCenter Server** - lists vCenter Server-related services if the current host has vCenter Server installed
- **Microsoft Exchange Server** - lists Microsoft Exchange Server-related services if the current host has Microsoft Exchange Server installed
- **Microsoft Dynamics CRM Server** - lists Microsoft Dynamics CRM Server-related services if the current host has Microsoft Dynamics CRM Server installed

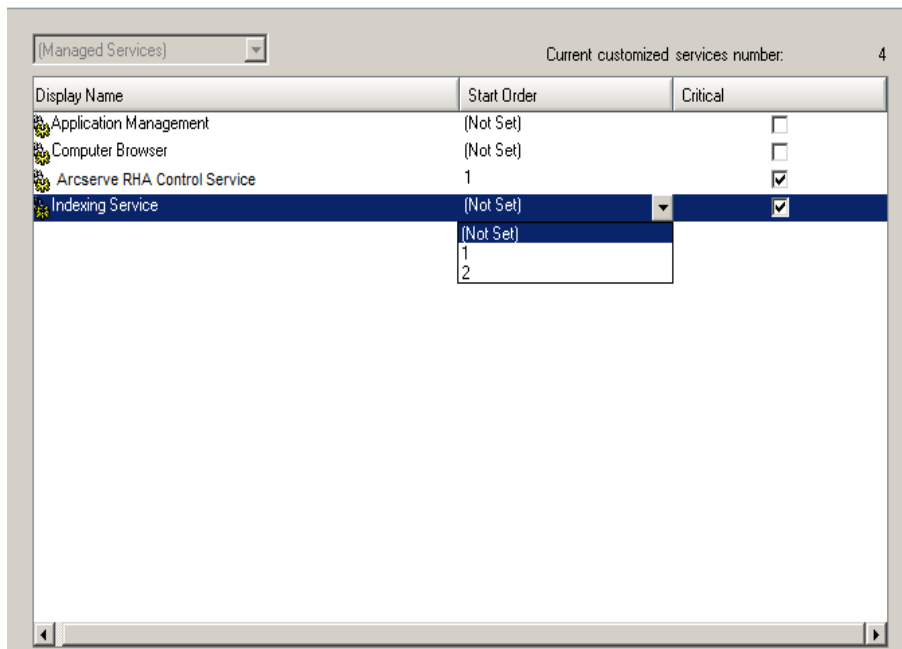
1. Select a Service to Monitor. Click the box to the left of each service listed to select it for monitoring.

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

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

Services Setting

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



3. In the Start Order column for each service you chose, specify the numeric value representing start order. For services where order does not matter, use the default value, (Not Set). The options available in the drop down list update

as you configure the value. The first service has only two options: Not Set and 1. The second service has three options: Not Set, 1 and 2, and so on. If you assign the same start order to two services, Arcserve Continuous Availability 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.

Run the Scenario from Outside the Wizard

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

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

To run the scenario from outside the wizard

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

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. For Exchange scenarios,


you should enable **Block Synchronization** and disable **Ignore same size/time files**.

5. Click the **OK** button. Synchronization may take a while, depending on data-base size and network bandwidth between the Master and Replica. You will receive the following message in the event window when the synchronization is complete: *All modifications during synchronization are replicated.*

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

Stop a Scenario

To stop a scenario

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

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

View a Report

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

The default storage directory of the generated reports is:

[ProgramFilesFolder]\CA\ArcserveRHA\Manager\reports

To view reports

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

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

The Report Center opens in a new window.

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

Available Reports per Scenario

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

Reports

Drag a column header here to group by that column

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

The Report Center consists of two tables:

- ◆ The upper table - Available Reports per Scenario - contains a list of all scenarios that have reports, with the type and number of available reports for each scenario.

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

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

The report you selected appears in a new window.

Synchronization 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 5: Switching Over and Switching Back

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

This section contains the following topics:

How Switchover and Switchback Work	42
Initiate Switchover	44
Initiate Switchback	45
Switchover Considerations	47

How Switchover and Switchback Work

After the Continuous Availability 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 Continuous Availability scenario configuration, Arcserve Continuous Availability sends you an alert or automatically initiates a switchover.

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

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

When you created an Continuous Availability 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.

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

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

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

Initiate Switchback

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

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

To initiate manual switchback

1. Ensure that both Master and Replica servers are available on the network and that the Engine is running.
2. Open the Manager and select the desired scenario from the Scenario pane.
3. Perform one of the following:

- ◆ If the scenario is already running, skip directly to Step 4
- ◆ If the scenario is not running, perform these steps and then go to Step 4:

- a. Click Run on the toolbar to start the scenario.

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

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

Note: See the *Arcserve Continuous Availability 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.

Click Perform Switchover on the toolbar while the scenario is running to reverse the server roles. A confirmation message appears.

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

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

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

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

Switchover Considerations

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

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

Chapter 6: Recovering Data

This section contains the following topics:

The Data Recovery Process	50
Setting Bookmarks	51
Data Rewind	52
Recover Lost Data from Replica	56

The Data Recovery Process

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

Arcserve Continuous Availability 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.

Setting Bookmarks

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

Notes:

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

To set a bookmark

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

The Rewind Bookmark dialog opens.

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

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

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

The bookmark is set.

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

Data Rewind

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

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

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

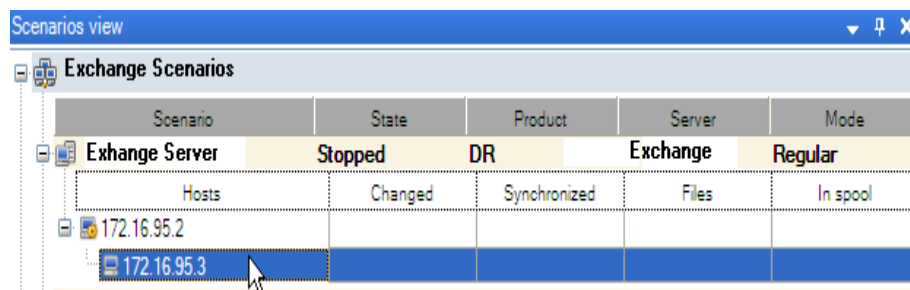
Important! The data rewind process operates in one way only - there is no replay forward. After rewind, all data subsequent to the rewind point will be lost, since data after the rewind point will be overwritten with new data.


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

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



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

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

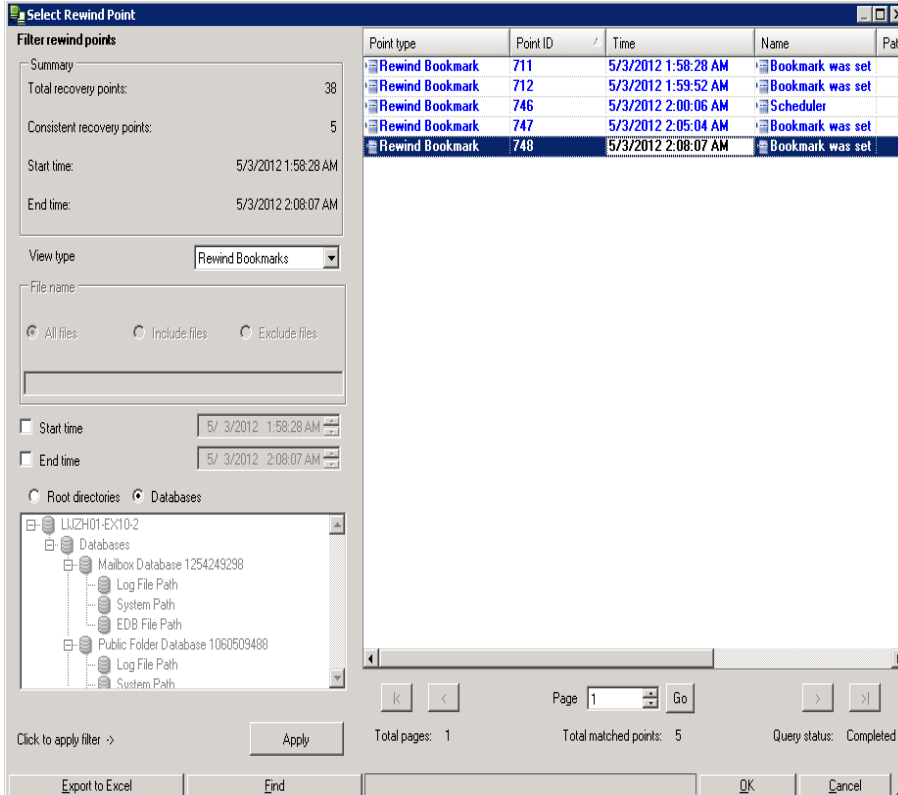
- Select one of the Rewind data options, depending on whether you want the rewind data synchronized back to the Master (option 2) or left on the Replica only (option 3).

Note: If the user credentials you used to log in to the Manager are different than the ones required for working with the Engine on the Replica, a **User credentials** dialog appears, asking you to enter log on account details for the selected Replica. For Exchange Server scenarios, option 3 is not needed.

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

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

The **Select Rewind Point** dialog appears.



Point type	Point ID	Time	Name	Path
Rewind Bookmark	711	5/3/2012 1:58:28 AM	Bookmark was set	
Rewind Bookmark	712	5/3/2012 1:59:52 AM	Bookmark was set	
Rewind Bookmark	746	5/3/2012 2:00:06 AM	Scheduler	
Rewind Bookmark	747	5/3/2012 2:05:04 AM	Bookmark was set	
Rewind Bookmark	748	5/3/2012 2:08:07 AM	Bookmark was set	

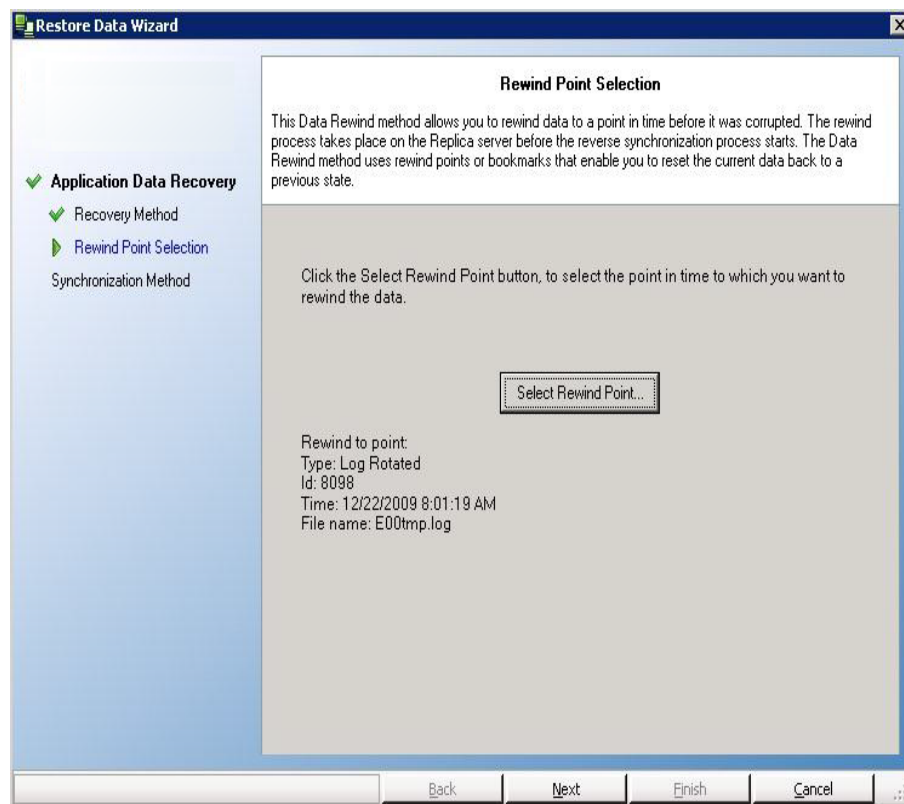
The **Select Rewind Point** dialog displays a list of all rewind points. These include modifications of folders and files that were automatically registered by the system and user-defined bookmarks.

Note: If the **Select Rewind Points** dialog is empty, make sure that the **Data Rewind** property is enabled.

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

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

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



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

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

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

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

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

Now, the Replication process can restart on the original scenario.

[Report Center Home Page](#)

Updated: Thursday, May 03, 2012 6:48:14 PM

Available Reports per Scenario							
SQL							
Exchange							
Scenario Name	Synchronization	Difference	Replication	Assessment Mode	Assured Recovery	Total Reports	
Exchange-DR	1	0	0	0	0	1	
Exchange - DR	3	0	41	0	0	44	
Exchange - HA	13	0	0	0	1	14	
Backward Exchange - HA	8	0	0	0	0	8	
Exchange-HA	5	0	0	0	0	5	
Backward Exchange-HA	1	0	0	0	0	1	
Exchange-MoveIP	3	0	0	0	0	3	
Backward Exchange-MoveIP	1	0	0	0	0	1	

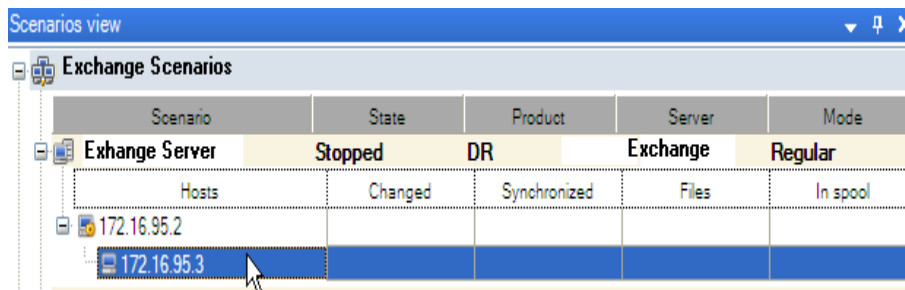
Recover Lost Data from Replica

In the following steps, File Server scenario screens are used as examples, but the procedures are similar for all server types.

To recover all lost data from a Replica

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

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



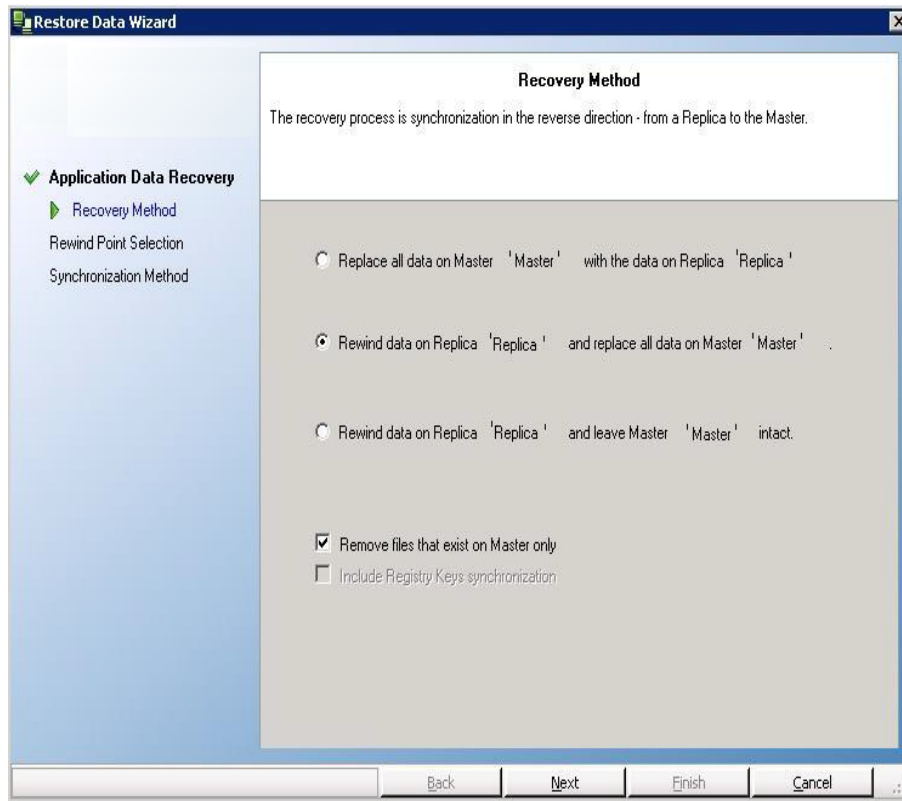
The Restore Data option is enabled.

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



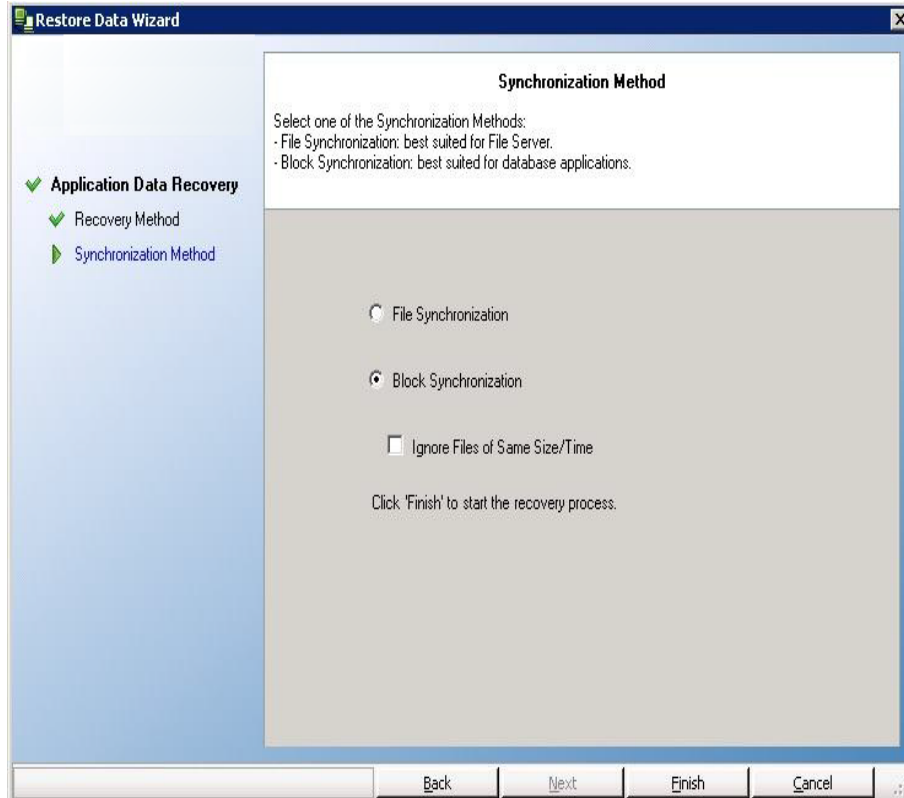
Note: If the user credentials you used to log in to the Manager are different than the ones required for working with the Engine on the Replica, a **User credentials** dialog appears, asking you to enter log on account details for the selected Replica.

The **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 (Block Synchronization). For more details, see the Arcserve Continuous Availability Administration Guide. Click **Finish**.

Once you finished initiating the recovery process, Arcserve Continuous Availability 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 reappears 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:

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
-------	-------	------------	-----------

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

Chapter 7: Redirection Methods

This section contains the following topics:

How Redirection Works	62
DNS Redirection	63
Move IP Redirection	64
Scripts Redirection	74
Mailbox Redirection	75
Disable Redirection	76

How Redirection Works

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

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

DNS Redirection

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

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

Using the record that has the Master server's name is the default, however you can configure Arcserve Continuous Availability 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](#).

This section contains the following topics:

- [Add IP on the Master Server](#)
- [Cluster Move IP](#)

Add IP on the Master Server

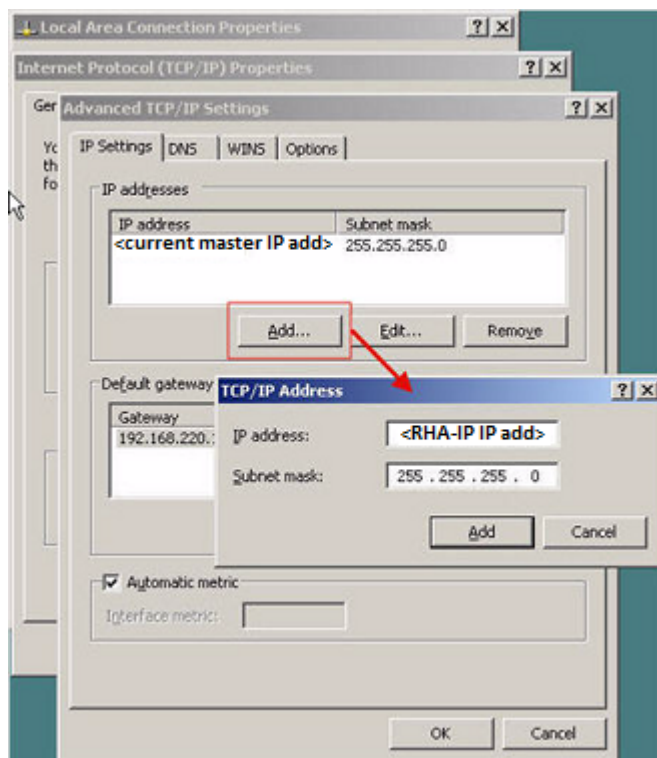
You need to add an additional IP address to the Master host, (which is denoted as *Arcserve-IP* in the following steps) to use Move IP redirection in your HA scenarios. This new IP address is used for Arcserve Continuous Availability internal communication and replication. This is necessary because once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

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

To add IP Address to Master Server

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

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



6. Click Add.

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

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

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

The procedures for both ways follow.

This section contains the following topics:

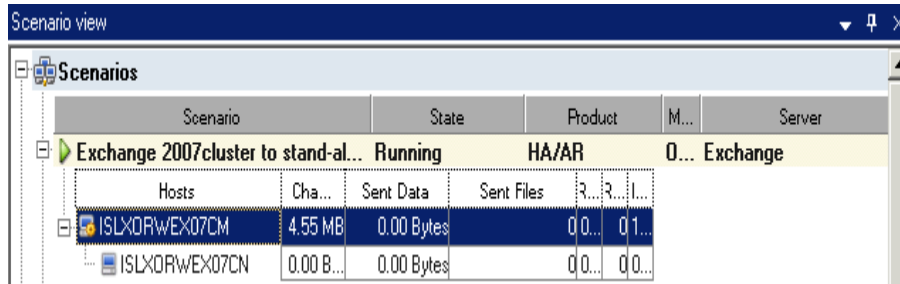
- [Add Arcserve-IP to Existing Scenarios](#)
- [Add Arcserve-IP to New Scenarios](#)

Add Arcserve-IP to Existing Scenarios

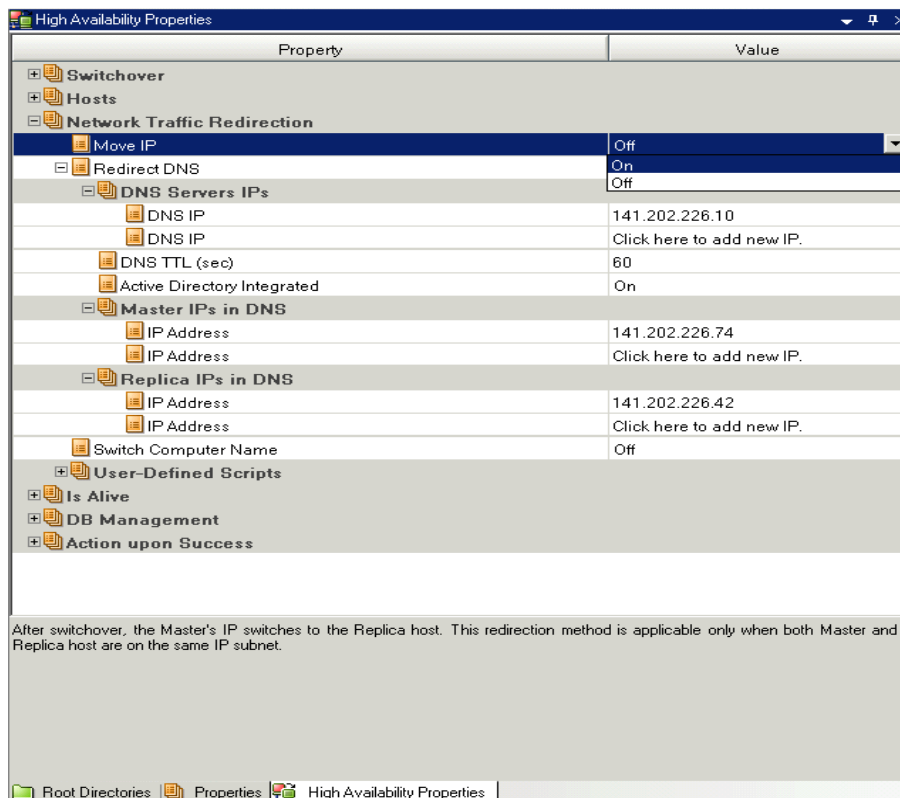
Perform this procedure only if you are using the Move IP redirection method.

To add the Arcserve-IP to existing scenarios:

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



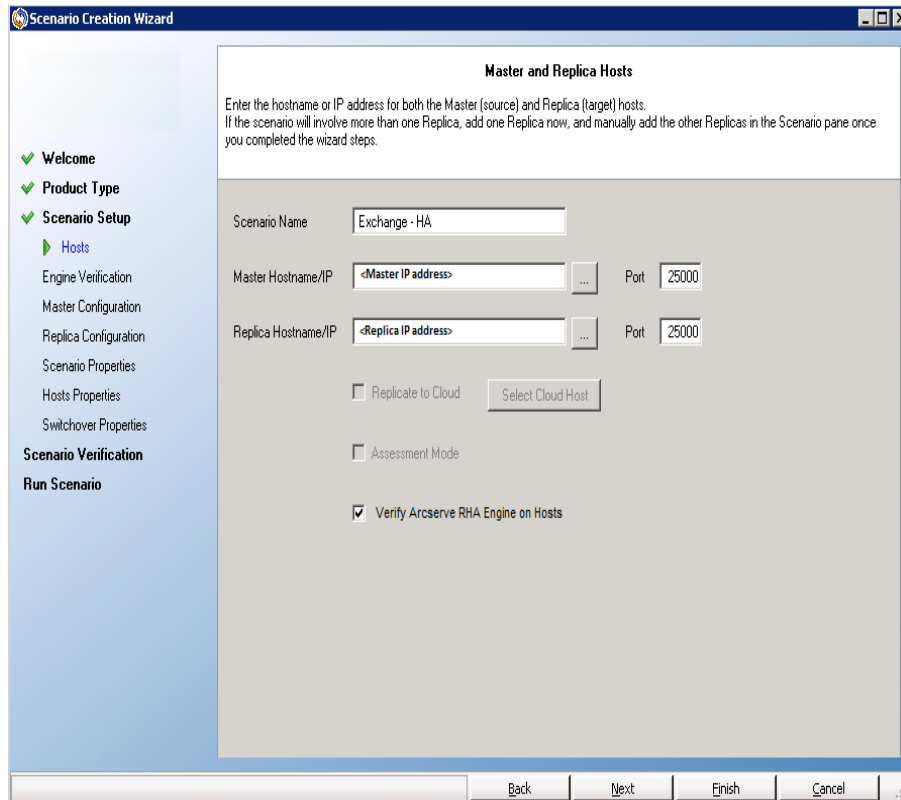
2. Right-click the Master and select Rename from the pop-up menu. Then, enter the Arcserve-IP address.
3. On the Framework pane, select the Switchover tab and then select the Replica server as the switchover host.
4. Set the Move IP option to On. Ensure that the IP address under Move IP, IP/Mask matches the production server IP address: this is the IP address that will switch over. If you are moving more than one IP address you can add multiple production IP addresses by selecting Click here to add new IP/Mask.



Add Arcserve-IP to New Scenarios

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

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



Cluster Move IP

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

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

This section contains the following topics:

- [Use the Master Cluster](#)
- [Use the Manager](#)
- [For New Scenarios](#)
- [For Existing Scenarios](#)

Use the Master Cluster

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

To use Cluster Move IP through the Master cluster

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

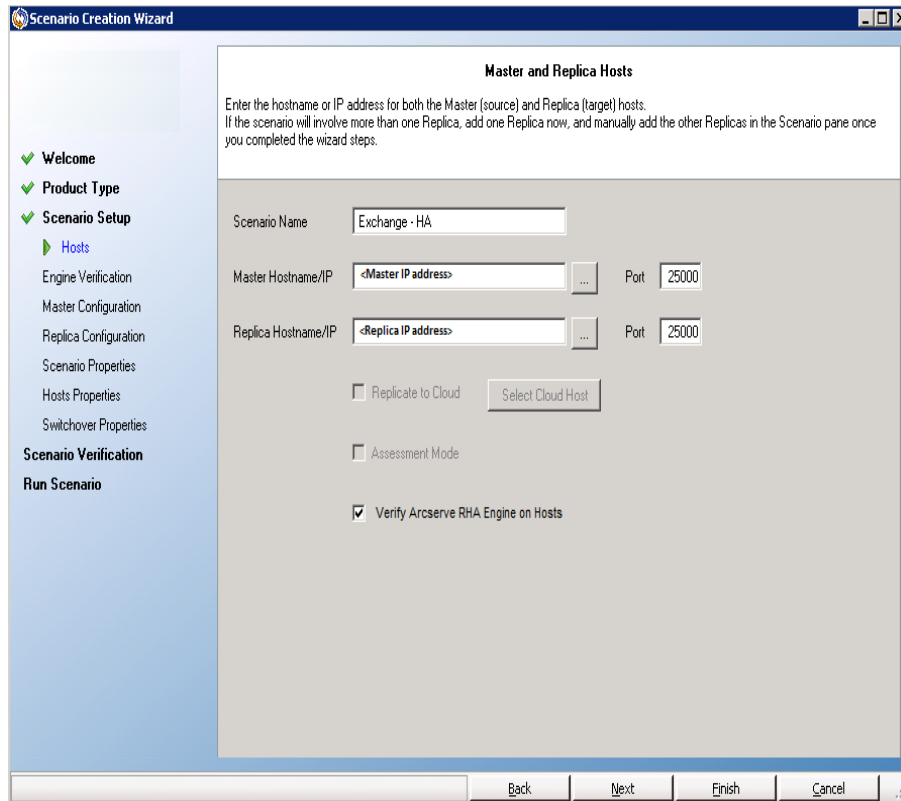
This new IP address is used for Arcserve Continuous Availability internal communication and replication. This is necessary since the current production IP address is not available on the Master cluster after switchover—it switches to the Replica server.

Use the Manager

This section details Cluster Move IP redirection using the Manager.

For New Scenarios


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



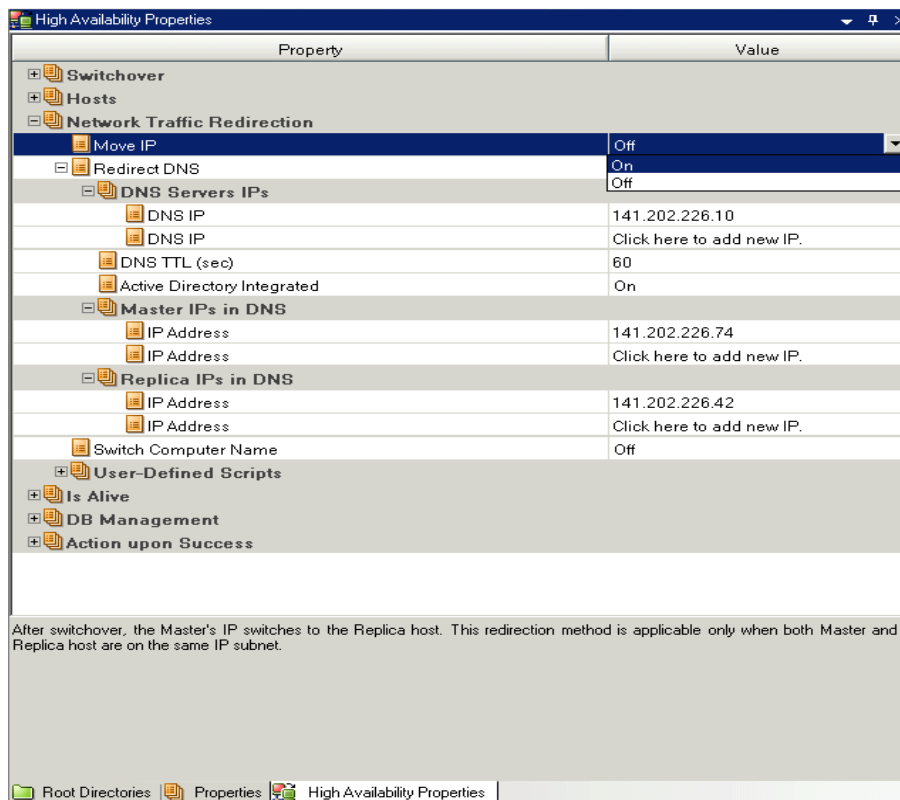
For Existing Scenarios

To use Cluster Move IP with existing scenarios

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



2. Right-click the Master and select Rename from the pop-up menu. Then, enter the RHA-IP address.
3. On the Framework pane, select the 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 the Master will switch over to. 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.

Scripts Redirection

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

Mailbox Redirection

To complete Switchover, all users mailboxes are redirected from the mailbox store on the Master to that on the Replica server. Continuous Availability redirects mailboxes with the following PowerShell commands:

- For Exchange Server 2007: Move Mailbox Store with the -ConfigurationOnly flag.
- For Exchange Server 2010: Set Mailbox with the -Database flag.

Disable Redirection

Affected Systems:

- Microsoft Exchange Server 2010

Microsoft redesigned the way in which Exchange Server manages traffic. The following ws_rep.cfg file property is used to control network traffic redirection on these Exchange systems.

Property	Value	Purpose
ExDisableRedirectDNS	True (default setting) False	<p>When set to True, disables DNS redirection and skips SPN redirection in the event a switchover occurs. Switchover itself occurs as expected.</p> <p>When set to False, the Exchange Management Console on the new active server is not started after switchover.</p> <p>When set to True, the Redirect DNS property is hidden in the Network Traffic Redirection group on the Arcserve Continuous Availability Manager screen.</p> <p>Set this value to True for Exchange Server 2010 HA scenarios.</p>

- Microsoft Exchange Server 2013

When you have not configured the Network Load Balancing and the CAS role is used for the master, then you have to enable the Redirect DNS option. To enable the Redirect DNS option, set the Redirect DNS option to on from the Network Traffic Redirection property in the High Availability Properties dialog.

Chapter 8: Additional Information and Tips

This section contains the following topics:

Spool Directory Settings	77
Recover Active Server	77
Recovering Servers	79

Spool Directory Settings

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

Place the spool folder on a drive with relatively low use such as a dedicated volume or boot/system volume. Do not place the spool folder on a volume containing frequently accessed system (OS), user, or application data. Examples include volumes containing databases, shared files, or the system pagefile. By default, the spool folder is located in the tmp folder under the Arcserve Continuous Availability 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.

Recover Active Server

In certain circumstances, it may be necessary to forcibly make the master or replica server the active server without completing the data synchronization process. For example, if switchover occurred but no data was changed on the replica server. In this case you may even have newer data on the master server making it

undesirable to synchronize data from the replica to the master server. Arcserve Continuous Availability allows for this option through a process called Recover Active Server. To use this option, ensure that the scenario is stopped, and select *Recover Active Server* from the Tools menu.

Important! While this option is the right choice in many situations, use it with caution. If used improperly data loss can occur. Normally, Arcserve Continuous Availability will not allow switchover from one host to another until all data is synchronized. It is designed this way so users are not redirected to an out of date data set that then overwrites what may be a more current data set. When using Recover Active Server, Arcserve Continuous Availability 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.

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

Recovering Servers

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

- First, perform the Recover Active Server procedure. For more information, refer to the topic, Recover Active Server.
- If the Recover Active Server procedure does not resolve the issue, try manually removing the IP address. For more information, refer to the topic, [Manually Recover a Failed Server when IP Redirection](#) is used.

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.

