

Microsoft IIS Server Operation Guide

Arcserve® Continuous Availability

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 is proprietary information of Arcserve and may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of Arcserve.

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® High Availability (HA)
- Arcserve® Replication
- 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
About This Guide	8
Related Documentation	9
Server Requirements	10
Base Configuration	11
IIS Server Configuration Requirements	12
Install IIS	13
Log On Account Conditions	14
Servers Operating in a Workgroup	15
About Clusters	16
IIS Metabase Replication	17
Register Arcserve Continuous Availability Licenses	18
Chapter 2: Creating and Using Scenarios	21
Create a New IIS Scenario	22
Scenario Properties	25
Manage Services	27
Run the Scenario from Outside the Wizard	30
Stop a Scenario	32
View a Report	33
Chapter 3: Redirection Methods	35
How Redirection Works	36
DNS Redirection	37
Move IP Redirection	38
Add IP on the Master Server	39
Add Arcserve-IP to Existing Scenarios	41
Add Arcserve-IP to New Scenarios	42
Cluster Move IP	43
Use the Master Cluster	44
Use the Manager	45
For New Scenarios	46
For Existing Scenarios	47
Switch Computer Name Redirection	48
Automatic Redirection Using Switch Computer Name	49
Scripts Redirection	50

Chapter 4: Switching Over and Switching Back	51
How Switchover and Switchback Work	52
Initiate Switchover	54
Initiate Switchback	55
Switchover Considerations	57
Chapter 5: Recovering Data	59
The Data Recovery Process	60
Set a Bookmark	61
Data Rewind	62
Rewind Data	63
Select Rewind Point Dialog	66
Recover Lost Data from Replica	68
Recover Active Server	70
Chapter 6: Additional Information and Tips	71
Spool Directory Settings	72
Configure Service Accounts	73
Additional Resources	74
Recovering Servers	75
Manually Recover a Failed Server-Move IP Address	76
Manually Recover a Failed Server-Switch Computer Name	77
Manually Recover Failed Server-IP and Switch Computer Name	78
Chapter 7: Install IIS 6.0 Management Compatibility for IIS 7.0/7.5	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:

About This Guide	8
Related Documentation	9
Server Requirements	10
Register Arcserve Continuous Availability Licenses	18

About This Guide

This document describes how you can implement a Arcserve RHA solution for Microsoft IIS 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](#)

Server Requirements

To implement Arcserve HA, refer to the appropriate list of requirements below, depending on the server type you selected. These components are licensed separately. If you do not have the license required to access support for a specific server type, contact Tech Support.

This section contains the following topics:

- [Base Configuration](#)
- [IIS Server Configuration Requirements](#)
- [Log On Account Conditions](#)
- [Servers Operating in a Workgroup](#)
- [About Clusters](#)
- [IIS Metabase Replication](#)

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 reside in the same Active Directory forest and are members of the same domain or a trusted domains.
- Both Master and Replica servers have Microsoft IIS installed.

IIS Server Configuration Requirements

- An instance of Microsoft IIS Server 6 or 7/7.5 installed on each server.
 - Both servers have the same IIS services installed-WWW, SMTP, and so on.
 - Both servers have identical web service extensions installed.
 - Full paths containing site files are identical on both servers.
- The standby server holds a clean installation of IIS only with the default sites.
- Sites on the Master server do not use URL redirection or UNC path redirection.
- (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.
- If anonymous access is enabled and used, configure the following:
 - To keep permissions synchronized between the two servers, both IIS processes should use the same user account for anonymous user access. Create a new domain user account and configure both IIS servers to use it. For more information about how you can do this, see the Microsoft documentation.
 - Although it is not specified in the Microsoft documentation, you must edit the local (or domain) group policy to allow the user account the following privileges: Allow log on locally, Allow log on as a batch job, and Access this computer from the network. Also, make sure to duplicate any permissions changes made to the file system for the original anonymous user account to the newly assigned domain account as well.
- If you are using SSL encryption, see the Microsoft documentation for information about copying the proper certificate.
- If you are using IIS 7.0/7.5, you must have IIS 6.0 Management Compatibility installed. Since the IIS 6.0 Management Compatibility is disabled by default when installing IIS 7.0/7.5, you must enable this option while installing IIS 7.0/7.5. For more information, see [Installing IIS 6.0 Management Compatibility for IIS](#).

Install IIS

To implement a Arcserve Continuous Availability solution for Microsoft IIS Server, you must first install IIS on the master and replica servers. To know the pre-requisites and the procedure to install IIS, see the Microsoft documentation.

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.
- 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.

Servers Operating in a Workgroup

For servers in a workgroup, set the Arcserve Continuous Availability Engine service account to Local System (unless you have added any custom scripts that require higher level permissions). Servers in a workgroup can use Redirect DNS only with DNS servers that allow non-secure updates. You can use Move IP, switch computer name, and custom redirection scripts normally.

About Clusters

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

The only configuration that requires some preparation is the use of IP Move with a cluster. For more information about how you can use Move IP with clusters, please see [Cluster Move IP](#).

IIS Metabase Replication

Arcserve Replication has been enhanced so that the IIS metabase is fully protected. You no longer are required to first create websites on the Replica before running an IIS HA scenario. While HA scenarios are running, the metabase is periodically replicated.

Note: Metabase protection is supported only with IIS 6.0 and 7.0/7.5 in Arcserve Replication HA scenarios.

High availability scenarios can now monitor and restore IIS services in addition to sites using the Is-alive check. This enables Arcserve Replication to recognize that if one particular IIS service is down, the corresponding site is not necessarily down as well. The Is-alive check monitors a list of IIS services participating in an HA scenario, checking the status for each and restoring any service that is down. This ensures Arcserve Replication does not attempt to restart a site that is still running.

IIS auto-discovery in a Microsoft Cluster Server environment has also been enhanced so that all local sites, or sites containing a path not on the shared MSCS disk, are identified and automatically excluded from the scenario, so that you may run scenarios with default settings.

Note: RHA only protects the meta data of the application pools that are associated with selected websites, other application pools are not synchronized or replicated.

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 2: Creating and Using Scenarios

This section contains the following topics:

Create a New IIS Scenario	22
Scenario Properties	25
Manage Services	27
Run the Scenario from Outside the Wizard	30
Stop a Scenario	32
View a Report	33

Create a New IIS Scenario

Creating scenarios is covered in full detail in the Arcserve Continuous Availability Administration Guide. This section provides additional information specific to IIS High Availability scenarios. 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 Microsoft IIS HA scenario

1. Open the Arcserve Continuous Availability Manager, choose Scenario, New, or click the New Scenario button.

The Welcome dialog opens.

2. Select Create New Scenario and then click Next.

The Select Scenario Type dialog opens.

3. Select IIS, High Availability Scenario, and Replica Integrity Testing for Assured Recovery (optional).

The Master and Replica hosts dialog opens. For more information on Assured Recovery, see the Arcserve Continuous Availability Administration Guide.

4. Enter a name for the scenario, the hostname or IP address and Port number for the Master and Replica servers. If either server is a MSCS cluster, enter the cluster resource virtual server name or IP address. Make sure that the Verify Arcserve Continuous Availability Engine on Hosts option is enabled (default).

5. Click Next.

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

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

7. [Optional] Select Replicate new sites in the selected root directories.

If you add new sites to a root directory after creating the scenario, Arcserve Continuous Availability replicates it automatically. You can see the new sites in the IIS Manager during replication.

8. Click Next.

The Replica Configuration dialog opens. Arcserve Continuous Availability compares the Microsoft IIS configurations between the Master and Replica hosts and then configures the replica automatically.

9. [Optional] Select Skip configuration on replica.
10. Click Next. The Scenario Properties dialog opens.
11. 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.

The Master and Replica Properties dialog opens. For more information, see [Scenario Properties](#) or the Arcserve Continuous Availability Administration Guide.

12. Accept default settings or change settings. Set a schedule for Integrity Testing for Assured Recovery and then click Next.
13. Wait for the Switchover Properties dialog to retrieve information. Configure the desired redirection properties and click Next.

The Switchover and Reverse Replication Initiation dialog opens. For more information, see [Switching Over and Switching Back](#).

14. Choose automatic or manual switchover, and automatic or manual reverse replication, as needed.

For more information, see Scenario Properties or the Arcserve Continuous Availability Administration Guide.

Note: Do not set both of these options to automatic.

15. Click Next.

Arcserve Continuous Availability starts verifying the scenario. You must resolve any errors reported and should resolve any warnings reported before continuing.

16. After the verification is successful, click Next.

The scenario is created.

17. 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](#).

-

Scenario Properties

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

The Properties pane and its tabs are context-sensitive and change whenever you select a different node from a scenario folder. You must stop a scenario before configuring its properties. Certain values cannot be modified once set; they are noted. For full details on configuring scenario properties and their descriptions, see the Arcserve 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 wish to change properties, and then select the appropriate tab. The following screen shows an IIS scenario as an example:

Settings on the Root Directories tab

Select a Master Server from the Scenario Pane. Double-click its Directories folder to add or remove Master Root Directories. Select or clear checkboxes next to folders, as desired, to include or exclude them. You may also edit directory names.

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

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, Scheduled, or periodic), 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
- ◆ 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 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 or choose email notification
- ◆ Report properties -- choose synchronization or replication reports, specify distribution or script execution
- ◆ (Replica) Scheduled Tasks -- set or suspend tasks, including Replica Integrity Testing for Assured Recovery. For more details, see the Arcserve Continuous Availability Administration Guide.
- ◆ (Replica) Recovery properties -- set delay or data rewind properties

Settings on the HA Properties Tab

These settings control how switchover and switchback are performed

- Switchover properties -- choose automatic or manual switchover, provide switchover hostname, and reverse replication settings
- Hosts properties -- specify the Master and Replica Fully Qualified Name
- Network Traffic Redirection properties -- choose Move IP, Redirect DNS, Switch Computer Name or User-defined scripts.
- Is Alive properties -- set the heartbeat frequency and check method
- DB Management properties -- instructs Arcserve 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

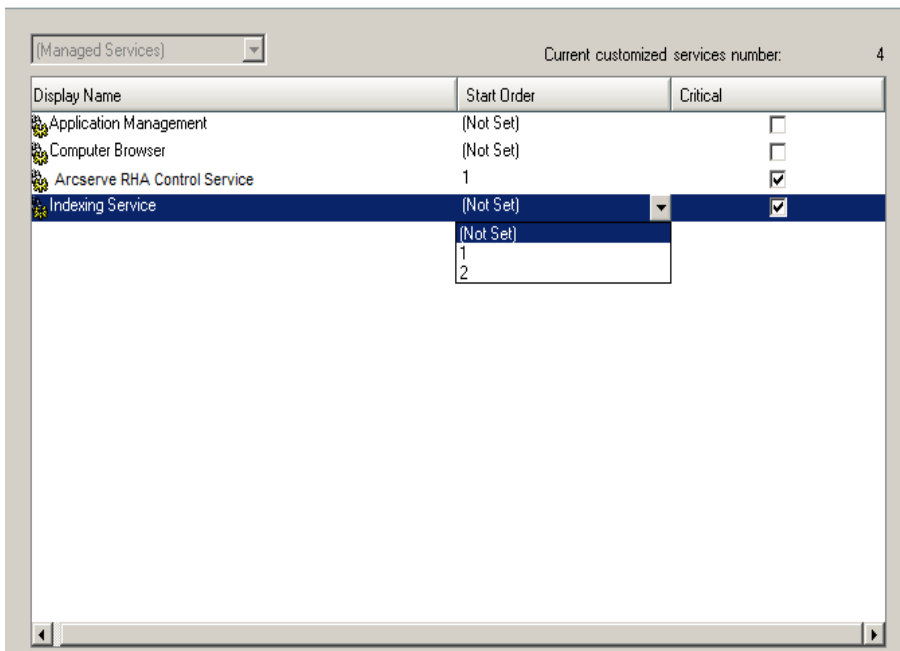
1. Select a Service to Monitor. Click the box to the left of each service listed to select it for monitoring.

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.

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 must restart the Engine on the Master before initiating replication.

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

To run the scenario from outside the wizard

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

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

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

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

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

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

4. Select 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.

5. Select Skip Synchronization only when you are certain that the files on both Master and Replica are identical.

Default: File Synchronization and Ignore files of same size/time are selected.

6. Click the **OK** button.

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

At this point, the scenario is operational and active. By default, a Synchronization Report is generated when synchronization finishes. To view the report, refer to the topic, [View a Report](#). You can also generate regular Replication Reports to monitor the replication process on each participating server. For more information, see the *Arcserve 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 3: Redirection Methods

This section contains the following topics:

How Redirection Works	36
DNS Redirection	37
Move IP Redirection	38
Switch Computer Name Redirection	48
Scripts Redirection	50

How Redirection Works

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

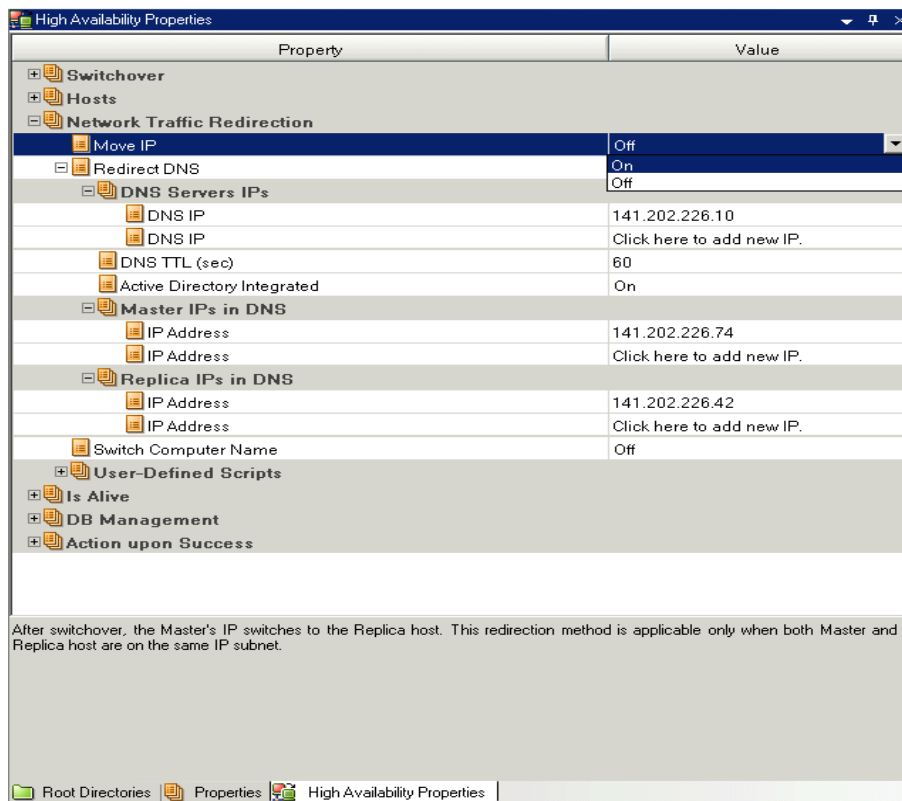
Perform this procedure only if you are using the Move IP redirection method.

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](#)

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](#)
- [For Existing Scenarios](#)

For New Scenarios

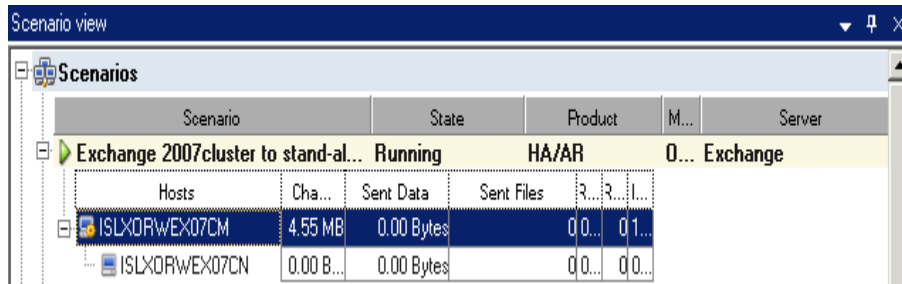
During the initial run of the Wizard, enter the RHA-IP and Replica IP addresses instead of the cluster virtual server names. The following illustration shows the RHA-IP entered in the Master Hostname/IP field and the Replica Server IP address entered in the Replica Hostname/IP field.



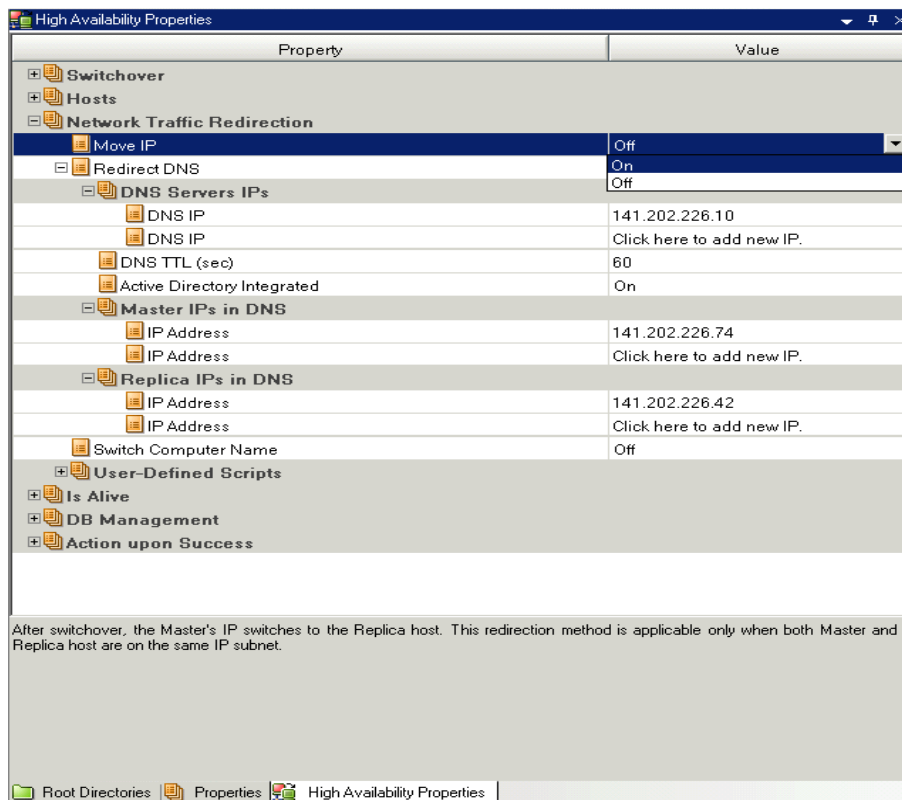
For Existing Scenarios

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.

Switch Computer Name Redirection

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

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

Arcserve Continuous Availability 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.

This section contains the following topics:

- [Automatic Redirection Using Switch Computer Name](#)

Automatic Redirection Using Switch Computer Name

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

Scripts Redirection

Arcserve 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.

Chapter 4: Switching Over and Switching Back

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

This section contains the following topics:

How Switchover and Switchback Work	52
Initiate Switchover	54
Initiate Switchback	55
Switchover Considerations	57

How Switchover and Switchback Work

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

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

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

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 Continuous Availability triggers Switchover without administrative involvement and could start Reverse Replication before you have investigated the cause of the failure. During Reverse Replication, Arcserve Continuous 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 5: Recovering Data

This section contains the following topics:

The Data Recovery Process	60
Set a Bookmark	61
Data Rewind	62
Select Rewind Point Dialog	66
Recover Lost Data from Replica	68
Recover Active Server	70

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.

Set a Bookmark

A *bookmark* is a checkpoint that is manually set to mark a state back to which you can revert. We recommend that you set 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.

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.

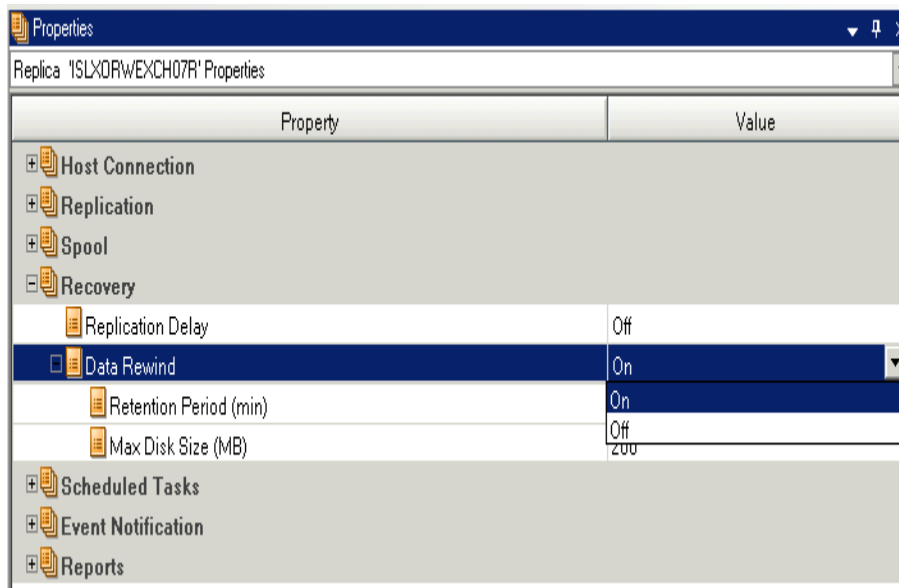
3. Accept the default name, or enter a meaningful name for the new bookmark, and click OK.

The bookmark is set.

Data Rewind

The Data Rewind recovery method lets you 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 to a previous state.

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



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

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.

More Information:

[Select Rewind Point Dialog](#)

This section contains the following topics:

- [Rewind Data](#)

Rewind Data

Important! The data rewind process operates in only one way—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.

To recover lost data using rewind points

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

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

4. Click Tools, Restore Data. If you are prompted for user credentials, enter the appropriate information and then 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 (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.

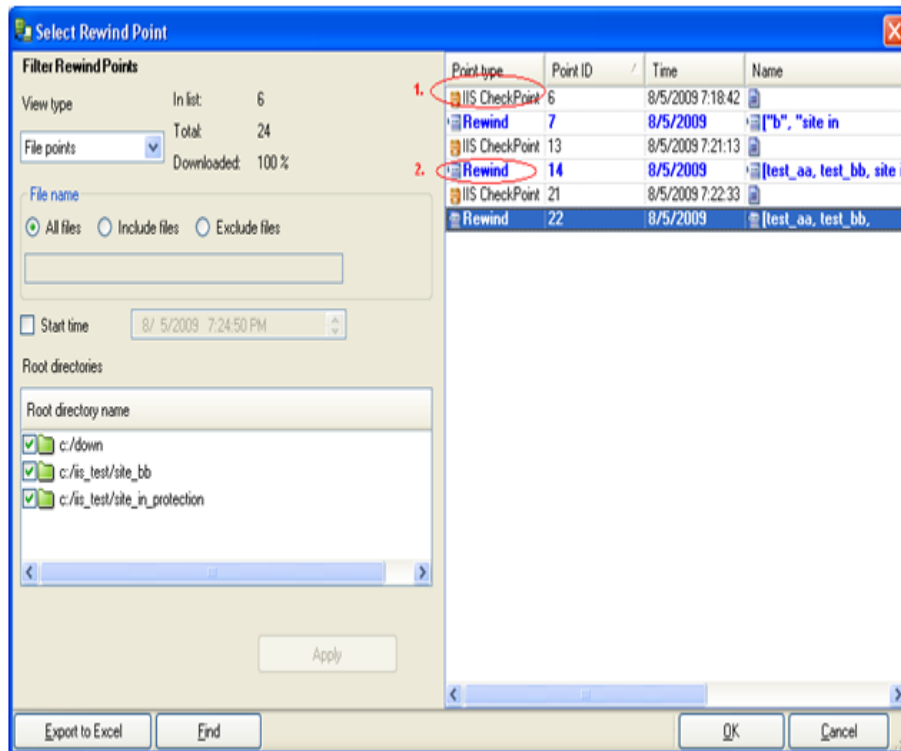
After you select a rewind data option, a recovery scenario is automatically created. This recovery scenario will run until the end of the rewind process.

6. Click Next.

The Rewind Point Selection window appears.

7. Click Select Rewind Point when it is enabled.

The Select Rewind Point dialog opens.



The Select Rewind Point dialog displays a list of all rewind points for IIS. These include modifications of folders and files that Arcserve Continuous Availability registered automatically and the bookmarks that you have defined.

You can filter the list 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.

The Rewind Point Selection window appears, with information about the rewind point you selected.

9. Click Next.

The Synchronization Method window appears.

10. Select the File Level Synchronization method and then 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 opens, 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 message, "Rewind process is completed successfully" in the Event pane.

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

The Synchronization Report is generated.

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

Select Rewind Point Dialog

The Select Rewind Point dialog displays a list of all the rewind points for IIS. These include modifications of folders and files that Arcserve Continuous Availability registered automatically and the bookmarks that you have defined.

The following sections are available on the left pane:

Filter Rewind Points

Lets you select a specific type of rewind point from the existing rewind points and displays statistics for it. The statistics are as given below:

- ◆ In list--The number of point types available in the selected list. These point types are displayed on the right pane.
- ◆ Total--The total number of points in the selected type of rewind point.
- ◆ Downloaded--The percentage of download of the selected point types.

File name

Indicates the types of files to be rewound. To include or exclude specific files, select the appropriate option button and enter the name in the box below.

Start time

Indicates the rewind process start time. Select the check box to change the rewind process start time.

Root directory name

Indicates the root directories that have been included in the rewind process. All the directories are selected; to exclude any directory, clear the check box and click Apply.

The following fields are displayed on the right pane:

Point type

Indicates the point type—IIS Checkpoint, file event, or user-defined bookmark. IIS Checkpoint is determined by the point when the metafile is closed after change; displaying the icon in this column indicates that some changes have occurred in IIS metadata. The metafile is updated on the Replica every 30 seconds. If you change any settings on the Master and immediately initiate a recovery process, you may not find the latest updates on the Replica.

Point ID

Indicates the point ID.

Time

Indicates the system time when the rewind happened.

Name

Indicates the name of the checkpoint.

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 dialog opens.

6. Make sure that the appropriate Synchronization method is selected. For more information, see the *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.

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

Recover Active Server

In certain circumstances, it may be necessary to forcibly make the Master or Replica server the active server without completing the data synchronization process. For example, if switchover occurred but no data was changed on the Replica server. In this case you may even have newer data on the Master server making it undesirable to synchronize data from the Replica to the Master server. Arcserve Continuous Availability allows for this option through a process called Recover Active Server.

To use the Recover Active Server option, stop the scenario, and then select *Recover Active Server* from the Tools menu.

Important! While this option is the right choice in many situations, use it with caution. If used improperly data loss can occur. Normally, Arcserve 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.

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

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

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

Chapter 6: Additional Information and Tips

This section contains the following topics:

Spool Directory Settings	72
Configure Service Accounts	73
Recovering Servers	75

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.

Configure Service Accounts

For servers in a workgroup, you must set the Arcserve Continuous Availability Engine service account to Local System (unless you have added any custom scripts that require higher level permissions).

To configure service accounts

1. Create identical user accounts on both servers.
2. Assign identical passwords to both accounts.
3. Configure this as the account used for anonymous access in the properties of each web site for which anonymous access is enabled.
4. If you want to use the default *IUSR_<computername>* account on both servers for anonymous access, ensure that you assign identical passwords to both accounts, and clear Allow IIS to control password in the Properties, Directory Security, Authentication Methods section of each web site for which anonymous access is enabled.

Important! If the IIS scenario is configured to replicate ACLs, a view of the security properties of the folder(s) containing web site content on the replica shows an unresolved SID – this is by design and can be ignored.

This section contains the following topics:

- [Additional Resources](#)

Additional Resources

The following Microsoft knowledge base articles discuss how to configure web site authentication:

- **For IIS 6.0:** *How To Configure IIS Web Site Authentication in Windows Server 2003* - <http://support.microsoft.com/kb/324274>

Recovering Servers

Arcserve Continuous Availability 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](#).
 - 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](#).
 - 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](#).

This section contains the following topics:

- [Manually Recover a Failed Server-Move IP Address](#)
- [Manually Recover a Failed Server-Switch Computer Name](#)
- [Manually Recover Failed Server-IP and Switch Computer Name](#)

Manually Recover a Failed Server-Move IP Address

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

To recover a failed server using Move IP Address redirection method

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

Manually Recover a Failed Server-Switch Computer Name

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

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

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

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

3. Connect to the network.
4. Rejoin the domain, make sure that you use the -RHA name assigned in step 2.
5. Reboot the computer.
6. If it is not already running, start the scenario from the Manager. (If automatic reverse replication was set to On, the scenario runs in backward mode so that the Replica server is now active and the Master server is now standby.)
7. Wait for synchronization to complete. Perform a manual switchover to make the Master server active. It is recommended that you do so outside of normal business hours.

Manually Recover Failed Server-IP and Switch Computer Name

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

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

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

Chapter 7: Install IIS 6.0 Management Compatibility for IIS 7.0/7.5

This section describes the necessary steps for installing IIS 6.0 Management Compatibility for IIS 7.0/7.5. This procedure is required if you want to create an HA scenario for IIS 7.0/7.5.


Note: If you intend to create an HA scenario for IIS 7.0/7.5, you need to repeat this process on both the Master and the Replica hosts.

To install IIS 6.0 Management Compatibility

1. On the Master or Replica host, open the Server Manager and select the Roles option. Then, click the Add Roles button.

The first page of the Add Roles Wizard appears.

2. Click Next. The Select Server Roles window appears.



3. Select the Web Server (IIS) check box, and then click Next.

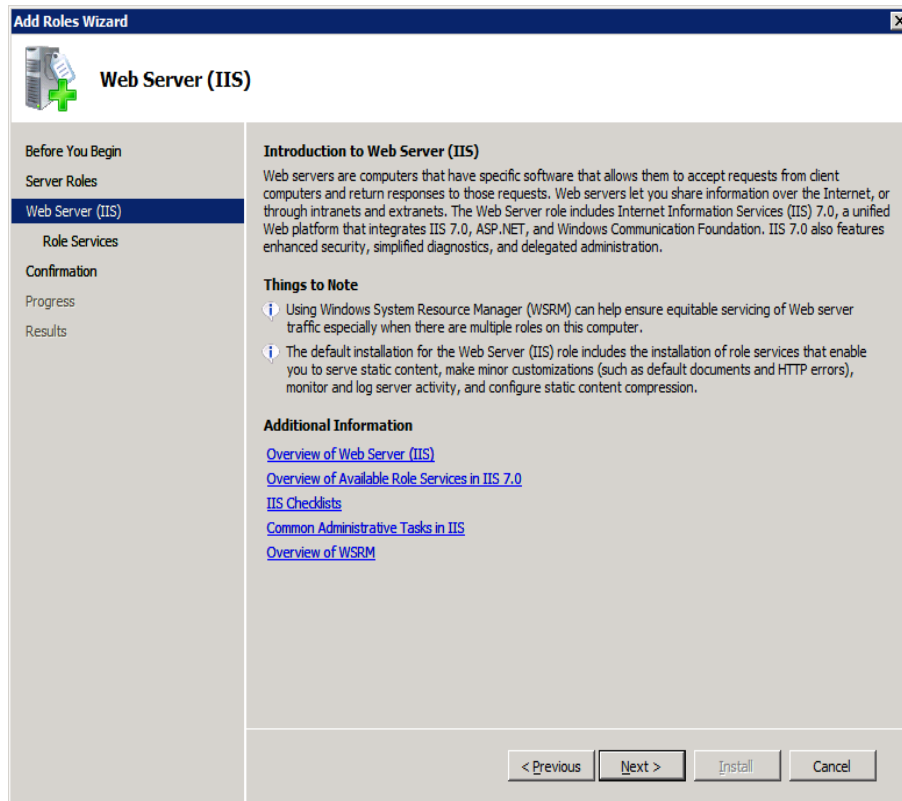
A pop-up message appears, asking you whether to add features required for Web Server (IIS).

4. Click Add Required Features.

The Select Server Roles window appears.

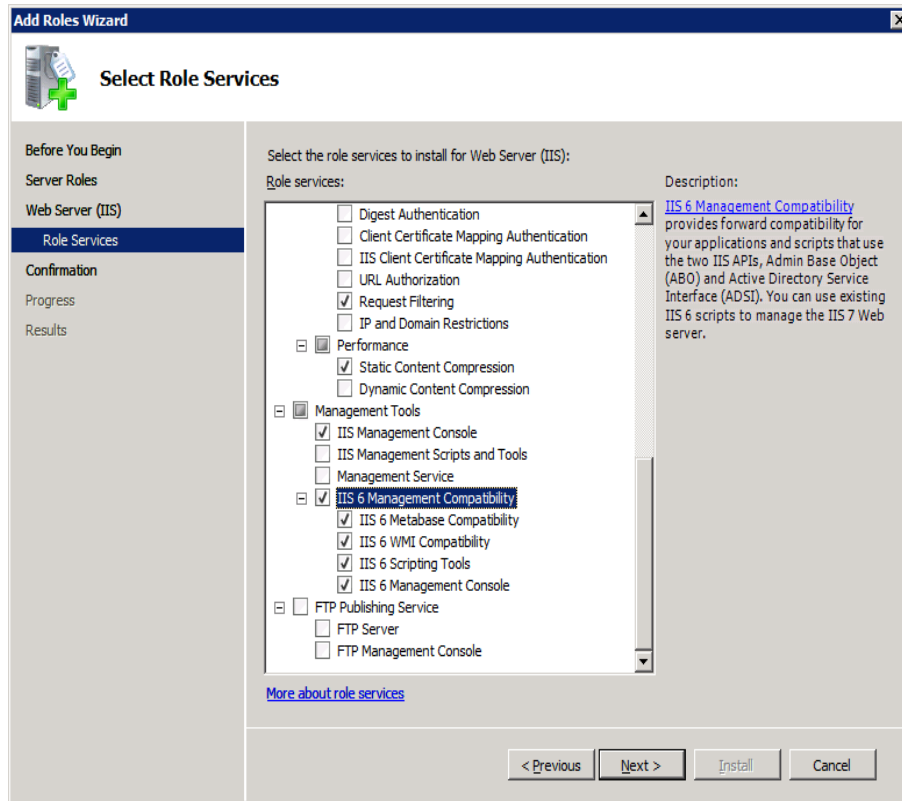
5. Click Next.

The Web Server (IIS) window appears.



6. Click Next.

The Select Role Service window appears.



7. On the Role Services list, select the IIS 6 Management Capability check box.
8. Click Next, and follow the Wizard instructions until the completion of the installation.

