

Arcserve Replication Integration Guide

Arcserve® Backup

19.0

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.

© 2022 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® Backup
- Arcserve® Unified Data Protection
- Arcserve® Unified Data Protection Agent for Windows
- Arcserve® Unified Data Protection Agent for Linux
- Arcserve® Replication and High Availability

Arcserve Backup Documentation

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

- [Arcserve Backup 19.0 Release Notes](#)
- [Arcserve Backup 19.0 Bookshelf](#)

Contact Arcserve Support

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

[Contact 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 that contains the 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 the 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.

Contents

Chapter 1: Integrating Arcserve Replication with Arcserve Backup	9
Introduction	10
Arcserve Backup	11
Arcserve Replication	12
Arcserve Backup and Arcserve Replication Integration	13
Capabilities and Benefits	14
Remote Branch Offices and Central Data Centers	15
Integration Terms and Definitions	16
How the Integrated Backup Process Works	20
Scenario Creation	21
Job Creation	22
Job Execution	23
How the Integrated Restore Process Works	25
Chapter 2: Installing and Configuring Arcserve Replication with Arcserve Backup	27
How to Install Arcserve Backup and Arcserve Replication	28
Integration Configurations	29
Configuration with a Stand-alone Arcserve Backup Server	30
Configuration with an Arcserve Backup Server Installed on a Replica Server-Arcserve Replication	31
Remote Branch Office (RBO) Configurations	32
RBO Configuration Example 1	33
RBO Configuration Example 2	34
RBO Configuration Example 3	35
Chapter 3: Performing Integrated Backup Jobs	37
Integrated Backup Jobs	38
Create an Arcserve Replication Scenario	39
Run an Arcserve Replication Scenario	42
Create and Run a Backup Job	45
Chapter 4: Performing Integrated Restore Jobs	49
Integrated Restore Jobs	50
Restore Using Arcserve Replication Failover	51
Restore Using Arcserve Replication Data Rewind	52
Restore Using Arcserve Backup	53
Set Arcserve Replication-specific Global Restore Option	54

Restore by Session	55
Restore by Tree	60
Restore by Query	62
Restore MS Exchange Scenario to Master Machine	64
Restore Microsoft SQL Server Database to Master Machine	65
Chapter 5: Monitoring Backup and Replication Jobs	67
Integrated Job Monitoring	68
Monitor Job Status Using Arcserve Backup	69
Job Queue Monitoring	70
Activity Log Monitoring	71
Monitor Job Status Using Arcserve Replication	72
Alert Notification	75
Arcserve Backup Alerts	76
Arcserve Replication Alerts	77
Report Generation	78
Arcserve Backup Reports	79
Arcserve Replication Reports	80
Chapter 6: Troubleshooting	83
Integrated Troubleshooting	84
Error and Warning Messages	85
Chapter 7: Glossary	87
assured recovery	87
continuous data protection	87
data rewind	87
failover	87
master server	87
replica server	88
synchronization	88

Chapter 1: Integrating Arcserve Replication with Arcserve Backup

This section contains the following topics:

Introduction	10
Arcserve Backup	11
Arcserve Replication	12
Arcserve Backup and Arcserve Replication Integration	13
Capabilities and Benefits	14
Remote Branch Offices and Central Data Centers	15
Integration Terms and Definitions	16
How the Integrated Backup Process Works	20
How the Integrated Restore Process Works	25

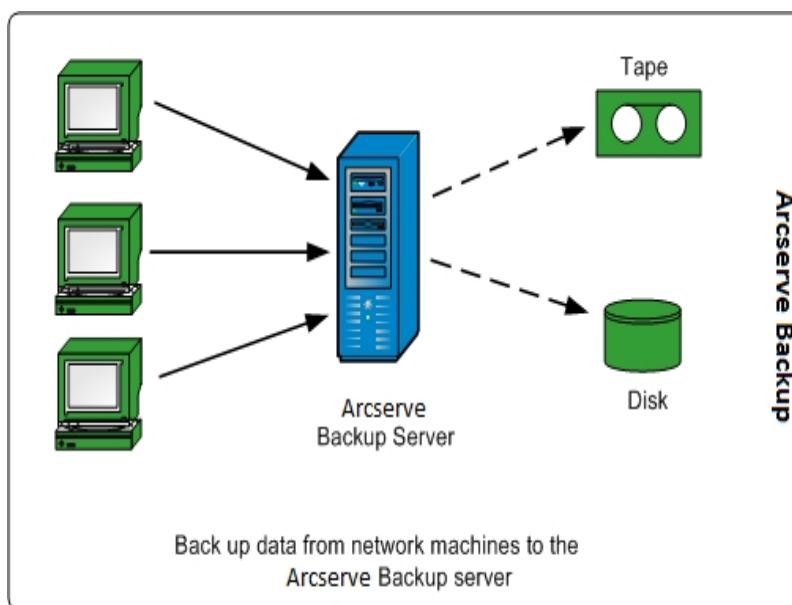
Introduction

Arcserve Backup provides you with high-performance disk-to-disk (D2D), disk-to-tape (D2T), disk-to-disk-to-tape (D2D2T), backup encryption and integrated anti-virus protection, multiplexing, and snapshot backup and recovery capabilities. The addition of Arcserve Replication complements these capabilities by adding continuous data protection, replication, and automated application failover. Together they provide a complete, 24-hour a day, 7-day a week integrated solution for recovery management, enabling you to better meet evolving compliance, business continuity, and disaster recovery objectives while saving time and resources.

Arcserve Backup

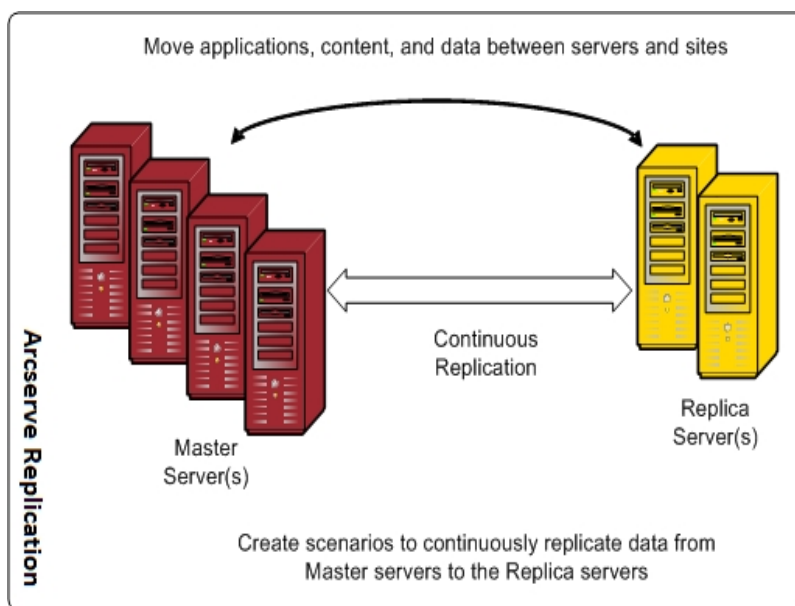
Arcserve Backup provides a complete, flexible, and integrated storage and recovery management solution for distributed and multiplatform environments. This application can back up and restore data from all the machines on your network (including machines running Windows, UNIX, and Linux) using optional client agents. Arcserve Backup also provides media and device management capabilities.

Arcserve Backup offers control from one management console and can support small-scale and large-scale enterprise environments consisting of one machine or many, across different platforms and organizations.



Arcserve Replication

Arcserve Replication is a data protection solution that uses asynchronous real-time replication to provide disaster recovery capabilities. This host-based software provides continuous data replication that transfers changes to application data as they occur to a standby replica server located locally or over the Wide Area Network (WAN). Continuous data replication ensures that the most recent data is always available for restoring purposes. Continuous data protection is based on Data Rewind technology for recovery from data corrupted by viruses, user error, or application error.

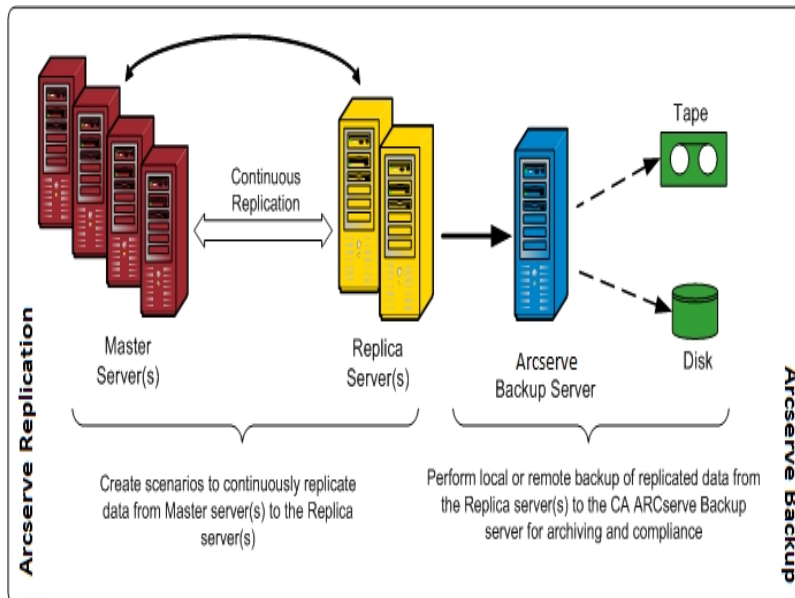


You can also improve your data protection capabilities by adding application monitoring and fully automated failover and failback. These capabilities are provided by Arcserve High Availability (HA), which is a high-availability solution for true continuous application availability.

Arcserve High Availability offers a superset of the capabilities of Arcserve Replication. In particular, Arcserve High Availability adds the capability for switch-over or automatic failover of clients from the production Master server to a secondary Replica server, without any need to reconfigure clients, as well as the ability to automatically monitor the status of the production server and the applications running on it. After the Master server is restored to its original state, Arcserve High Availability lets an IT administrator reinstate the Master server with the push of a button, automatically failing back from the Replica server with no loss of data or application availability.

Arcserve Backup and Arcserve Replication Integration

The integration between Arcserve Backup and Arcserve Replication provides the dual benefit of continuous data protection and backup of this protected data. Through this integration, Arcserve Replication continues to provide real-time, continuous data replication from the Master server (production server) to the Replica server, while Arcserve Backup then backs up this replicated data from the Replica server to the Arcserve Backup server for archiving and compliance. By performing the backup operation from the Replica server, no backup window is required and the impact on the Master server is minimized, enabling the Master server to continue working without any performance degradation. In addition, you can recover the backed-up data to the Master server or to the Replica server using Arcserve Backup.



Capabilities and Benefits

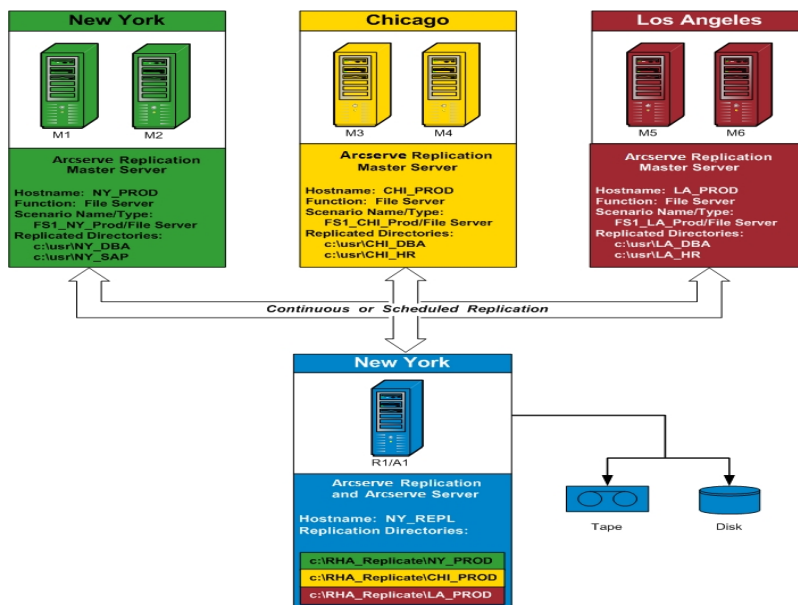
The integration between Arcserve Backup and Arcserve Replication lets you use replication and continuous protection (using Arcserve Replication) and backup (using Arcserve Backup) for archiving and compliance.

The integration provides the following core capabilities and benefits:

- Use of the familiar Arcserve Backup interface for backup job configuration, scheduling, and management.
- High-performance D2D, D2T, and D2D2T backup and snapshots.
- Backup encryption and integrated antivirus protection.
- Built-in device and media management.
- Real-time, continuous replication of files and data as they change to ensure that the most recent data is always available for restore or recovery.
- Backups are done from the Replica server, which minimizes the impact on the Master server and provides an unlimited backup window.
- Multiple recovery options as follows:
 - Recovery through Arcserve Replication using Failover. Recovery capability of an application to a standby server at another location and switch-over failback of the application once the production server is restored.
 - Recovery through Arcserve Replication using Data Rewind. Recovery capability to rewind to any previous point in time (rewind points) to recover from corrupted data, which provides continuous data protection (CDP).
 - Recovery through Arcserve Backup. Recovery capability from tape or other traditional backup media.
- Real-time consolidation of data from multiple remote branch offices (RBOs) to a central data center (CDC), which provides centralized backup and data consolidation. You can use the replica residing in the central data center to perform centralized backup and reduce the need for IT support at every location.
- Assured Recovery capability to provide nondisruptive, fully automated, in-depth testing of the disaster recovery replica server. This allows recoverability testing of the application on the Replica server without any disruption to the Master server, to the replication process, or to the automated failover protection mechanisms that are in place in case of a disaster.

Remote Branch Offices and Central Data Centers

Arcserve Backup together with Arcserve Replication lets you create backup replicas of remote branch office (RBO) servers at a central data center (CDC). The backup replica servers located at a CDC provide better disaster recovery and continuous application availability for your branch office servers, as well as centralized and consolidated backups of multiple branch offices at a single facility. This consolidation of backups of remote branch offices significantly reduces the need for competent and sophisticated IT support at every location. Installing and maintaining multiple servers, storage devices, and applications in many different locations can be expensive. In addition, the backup replica servers at a CDC also reduce the security risk associated with tape transport and offsets tape media and handling costs.



Integration Terms and Definitions

Before you can understand the details of the integration between Arcserve Backup and Arcserve Replication, you must be familiar with some of the terms and definitions used by each product.

The integration uses the following terms and definitions:

Scenario

An Arcserve Replication scenario is the fundamental basis for managing the operation of the system. An Arcserve Replication scenario always includes at least one Master server and one Replica server. In addition, multiple independent scenarios can run on a single server.

A scenario is a structure that describes the following:

- What applications and data are to be protected
- Where they are located (that is, the Master server and source directories)
- Where the data is to be replicated (the Replica servers and target directories on them)
- Whether and how automated failover and testing should take place

Note: Arcserve Backup supports File Server, MS Exchange, and SQL Server scenarios only.

Master Server

The Master server is the active or production server that lets you actively change (read and write) data. Any changes made at any given time on the Master server are captured continuously in real time and transferred (or replicated) to one or more of the associated Replica servers so that all the Replica servers contain an exact copy of the data on the Master server at all times.

Replica Server

The Replica server is the passive server. This is the server from which data cannot be changed (read only) in any way except through changes replicated from the Master server. There can be multiple Replica servers associated with a single Master server. When synchronizing the replicated data with the Master server, the data between the servers is compared and only the changes made to the Master server are sent to the Replica servers, which minimizes WAN traffic.

Continuous Data Protection (CDP)

Continuous data protection (CDP) is the ability to recover data not just to certain isolated previous states captured, for example, in a daily or weekly backup or snapshot, but to recover the data back to any point in time. That way, whenever a virus occurs, you can recover to a point just minutes before the virus occurred with essentially zero data loss and a very fast recovery time.

Failover

Failover is a feature that detects when the protected application on the Master server fails and switches (either automatically or manually activated) to a designated Replica server with essentially zero loss of data and time. If a Master server fails or must be shut down for maintenance, a synchronized replica (locally or in a remote location) will instantly and automatically take its place.

Data Rewind

Data rewind is a recovery method that allows rewinding files to a point in time (rewind point) before they were corrupted. This built-in rewind technology occurs on the Replica server and any "data rewinding" to a previous state can be performed on the Replica server only.

Rewind Point

A rewind point is a checkpoint in the Rewind log marking an event or operation. The actual information stored includes the operation that will undo the event in case the rewind point is activated. Data recovery uses these rewind points or bookmarks in the Rewind log to reset the current data back to a previous state.

Assured Recovery

Assured Recovery lets you perform a real test of your disaster recovery server by actually running the application, including modifying data, without impacting your production environment or your previously replicated data.

Using Assured Recovery, you can perform transparent, non-disruptive testing of a replicated data snapshot to start up application services and perform all operations necessary to verify the integrity of the data on the Replica server.

Assured Recovery provides this functionality without ever leaving your production systems vulnerable during testing, without disrupting production application availability in any way, and without having to resynchronize the data after testing is complete.

Suspend Mode

The Suspend mode temporarily stops delivering changes to the suspended Replica server. Changes will continue to be recorded in a spool until replication is resumed so that resynchronization is not required. After replication is resumed,

the accumulated changes are then transferred and applied without any need to perform a full resynchronization data.

When you back up a scenario with Assured Recovery configured, the backup is application-consistent and requires no application recovery after restore.

However, when you back up a scenario with the Suspend mode enabled (no Assured Recovery configured), the backup may require application recovery after restore, depending upon the state of the application at the time of the backup.

Synchronization

Synchronization is the process of bringing the data on the Replica server in sync with the data on the Master server. To properly synchronize the Master server and the Replica server, their two file structures are compared to determine what content (files and folders) on the Master server is missing or is different from that on the Replica server. The levels of synchronization are as follows:

- File-level synchronization involves replicating an entire file when a change occurs. This process, which is used for smaller files, copies the entire set of data and sends it to the Replica server (if no part of it exists on the Replica server).
- Block-level synchronization involves determining what has changed and sending only the changes to the Replica server (to minimize the bandwidth and time required). This process is used for replicating large data sets such as databases.

Replication

Replication is a process that maintains identical copies of files and databases by real-time capture of byte-level changes in files on the Master server. These captured changes are asynchronously transmitted to the Replica servers. Because replication continuously updates the source data to another machine, a Replica server always contains the same data as in the Master server. To avoid attempting to restore files that are in use, the application needs to be not running (off-line).

Entity

Defines the granular level of detail for an Arcserve Replication scenario for backup and restore purposes. The level of granularity for an entity depends on the type of scenario.

File Server Entity

For a File Server scenario, an entity represents all of the files and directories belonging to the same volume on a Master server.

For example, on a Master server the contents of the C: drive would constitute one entity, while the contents of the D: drive would be a separate entity.

SQL Server Entity

For a SQL Server scenario, an entity represents a SQL database.

For example, on a Master server the contents of the Company A Employees database would constitute one entity, while the contents of the Company B Employees database would be a separate entity.

MS Exchange Entity

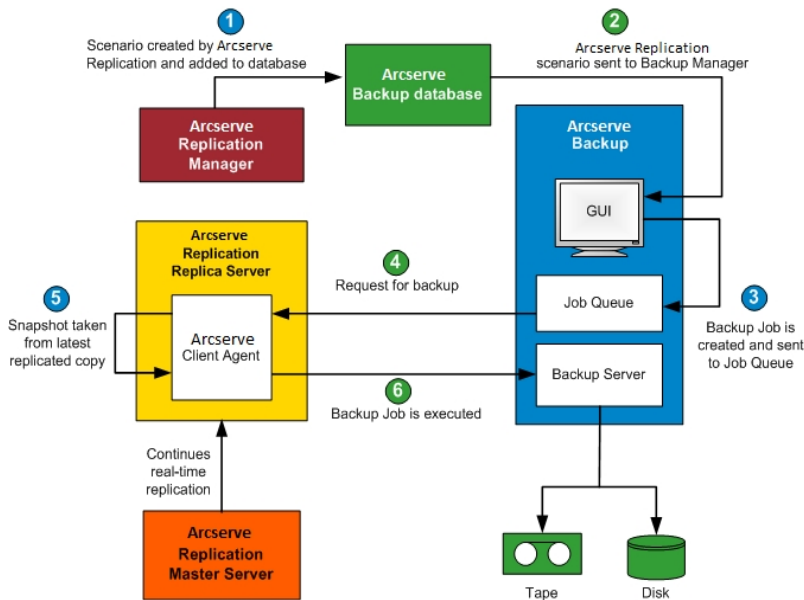
For an MS Exchange scenario, an entity represents an MS Exchange Storage Group. The entity name is the MS Exchange Storage Group name.

How the Integrated Backup Process Works

The backup process will back up everything that is part of a scenario. The backup job will create a session for each entity that constitutes a scenario. The supported scenarios are SQL Server, MS Exchange, and File Server.

The integrated backup process for Arcserve Backup and Arcserve Replication consists of the following three basic functions:

- [Scenario Creation](#)
- [Job Creation](#)
- [Job Execution](#)



Scenario Creation

From the Arcserve Replication Manager, a scenario is created and the related information is inserted into the Arcserve Backup database. For an existing scenario, the related information can be inserted into Arcserve Backup database by using the Update Arcserve Backup Server option, which is accessible from the Tools menu of the Arcserve Replication Manager. Arcserve Backup then queries the database, becomes aware that the scenario exists, and presents the scenario to the user through the Backup Manager GUI. The Arcserve Replication scenarios are listed in the source tab of the Backup Manager. When a scenario is selected from the Backup Manager, some corresponding properties are displayed to provide additional information about the selected scenario. These displayed properties include Scenario Name, Scenario Type, and other relevant information about the Master server and Replica servers.

Note: An Arcserve Replication scenario always includes at least one Master server and one Replica server. In addition, multiple independent scenarios can run on a single server.

You can remove a scenario from the Arcserve Backup database by using the Delete Machine/Object option, which is accessible from the properties pane of the Backup Manager GUI. Using Delete Machine/Object removes the scenario from the Arcserve Backup database only and not from the Arcserve Replication Manager.

More information:

[How the Integrated Backup Process Works](#)

[Job Creation](#)

[Job Execution](#)

Job Creation

From Arcserve Backup, a backup job is created with the user specifying the source, destination, and any other typical backup options. When you attempt to submit a backup job, you will be prompted to provide two sets of security credentials (user name and password). One set of credentials allows the Arcserve Backup server to access and communicate with the agent on the Replica server and the other set of credentials allows Arcserve Backup to log in to the Master server to prepare for the backup. After the required security credentials are entered, the backup job is submitted to the job queue to be executed at the scheduled time.

More information:

[How the Integrated Backup Process Works](#)

[Scenario Creation](#)

[Job Execution](#)

Job Execution

At the scheduled time, Arcserve Backup connects to the agent running on the Replica server and asks Arcserve Replication to create shadow copies of the volumes that are hosting the replicated data for the scenario being backed up. When the request is received, the continuous real-time replication of the scenario is temporarily suspended to facilitate the creation of the shadow copy. After the shadow copy is created, the Replica server resumes performing real-time replication and the agent running on the Replica server proceeds to perform the backup from the shadow copy. You can install the Arcserve Backup server on the Replica server or as a separate dedicated server.

For SQL Server and Exchange scenarios, when the backup is complete, the shadow copy is preserved and removed only if the maximum number of retained shadow copies is reached. By default, Arcserve Replication creates and preserves 10 shadow copies and then begins replacing the oldest shadow copies with newer ones. You can change the setting for the number of shadow copies preserved using the Number of Shadow Copies Set To Keep property on the Arcserve Replication Manager.

Note: For more information on setting the number of shadow copies to be kept, see the [Arcserve RHA Administration Guide](#).

For File Server scenarios, when the backup is complete, the shadow copy is preserved and not removed at all.

When you back up a scenario with Assured Recovery configured, the backup is application-consistent and requires no application recovery after restore. When you back up a scenario with Suspend enabled (no Assured Recovery configured), the backup may require application recovery after restore, depending upon the state of the application at the time of the backup.

Performing the backup on the Replica server allows the Master server to continue working without any interruptions or performance degradation. In addition, all backup catalog information is recorded as if the backup was performed on the Master server, ensuring that the restore view of the data will always be the same as if the backup was taken directly from the Master server. You can recover the backed-up data to the Master server or to the Replica server using Arcserve Backup.

In addition, for multistreaming backup jobs, each Arcserve Replication scenario is backed-up as a child job. If one node contains multiple scenarios, the master job will split them so that each child job will back up one scenario.

More information:

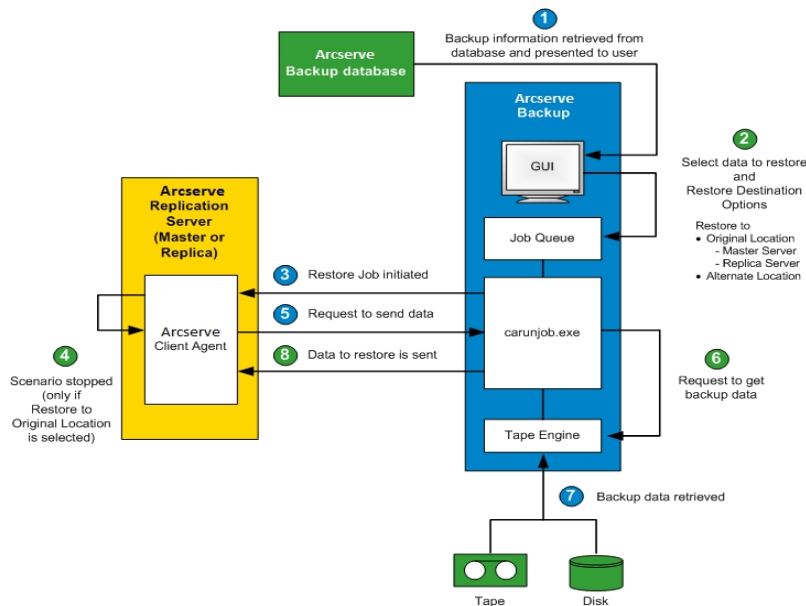
[How the Integrated Backup Process Works](#)

[Scenario Creation](#)

[Job Creation](#)

How the Integrated Restore Process Works

The restore process will recover data that was replicated from the Master server and then backed up using Arcserve Backup. For SQL Server scenarios and MS Exchange scenarios, the restore process supports only full scenario restores. For File Server scenarios, the restore process supports either full scenario restores or granular restores of files, directories, and volumes.



From Arcserve Backup, a list of machines that were backed up is retrieved from the database and displayed through the source tab of the Restore Manager GUI. When you select a source from the Restore Manager, some corresponding properties are also displayed to provide additional information about the selected source. After you select a source, you must also select the destination for the restore. The available destination options are to restore the data to its original location or to an alternate location. If you choose the restore to the original location option (default), you will then have to select whether to restore to a Master server or a Replica server for that location. If you do not choose the restore to original location option, you must browse to locate the alternate location for the destination.

When restoring to a Replica server, ensure that the corresponding application is not running to avoid attempting to restore files that are currently in use. When restoring to the Master server, ensure that the Arcserve Backup Client Agent is installed and running.

For SQL Server and MS Exchange, if the corresponding application is running and you restore an Arcserve Replication scenario to the original location, the restored files are temporarily created with a .TMP extension. After the restore job is

successful, you are prompted to reboot the server to overwrite and replace the existing active files. When the server is rebooted, the restored files are merged into the original database. If the corresponding application is not running, temporary files will not be created during the restore to original location process and you will not need to reboot the server.

Rebooting a SQL Server installed on a cluster environment is not feasible, so you must stop the application resource before performing a restore to the original location.

When a restore job is submitted, you are prompted to provide security credentials (user name and password) based on the following options:

- If you selected to restore to the original location, depending on the restore option specified, you must provide either one set of credentials to allow Arcserve Backup to communicate with the agent on the Master server (restore to master) or two sets of credentials to communicate with both the Master server and Replica server (restore to replica).
- If you selected to restore to an alternate location, you must provide only one set of credentials to login to the agent on the machine where the restore will be performed.

After you enter the required security credentials, the restore job is submitted to the job queue for execution at the scheduled time.

At the scheduled time, Arcserve Backup connects to the agent running on the Arcserve Replication server (master or replica, as specified) to initiate the request for data. If you selected the restore to original location option, the continuous replication of the scenario is temporarily suspended before the agent requests the data. Suspending the scenario replication helps ensure data consistency between the Master server and the Replica server. The data on the Master server could be different from the restored data on the Replica server and possibly corrupted. In this case, if the continuous replication process was not suspended, the newly restored data on the Replica server would then be overwritten with the corrupted data from the Master server. If the scenario replication cannot be stopped, the restore job will fail. If necessary, you can select a Restore Manager Global Option to continue the restore job even when the scenario cannot be stopped.

When the agent running on the Arcserve Replication server contacts Arcserve Backup and requests the data to be sent, the data is retrieved from wherever it was stored (tape or disk) and sent to the specified destination. This process of requesting and sending data is repeated until all the data from the backup is restored. After the restoration is complete, you can restore the data on the Replica server back to the Master server.

Chapter 2: Installing and Configuring Arcserve Replication with Arcserve Backup

This section contains the following topics:

How to Install Arcserve Backup and Arcserve Replication	28
Integration Configurations	29
Remote Branch Office (RBO) Configurations	32

How to Install Arcserve Backup and Arcserve Replication

The following process for installing Arcserve Backup and Arcserve Replication does not change for the integrated products:

- Perform the usual Arcserve Backup installation on the server designated as the Arcserve Backup server.

Note: For more information about installing Arcserve Backup, see the [Arcserve Backup Implementation Guide](#).

- Perform the usual Arcserve Replication installation on the servers designated as the Master server and Replica server.

To perform backups the following configurations must exist:

- The Replica server must be a Windows server configured according to the compatibility matrix posted for Arcserve Backup and Arcserve Replication.
- The Master server must be a Windows server configured according to the Arcserve Replication compatibility matrix.

Note: For more information about installing Arcserve Replication, see the [Arcserve RHA Administration Guide](#).

Integration Configurations

You can choose two integration setups, depending on where the Arcserve Backup server is installed. One setup configuration involves installing the Arcserve Backup server on the Replica server, while the other setup configuration involves having the Arcserve Backup server as a separate stand-alone server.

The following chart lists the supported configuration versions for integration between Arcserve Backup and Arcserve Replication:

Arcserve Backup	Arcserve Replication	Integration Supported?
18.0	r16.5 SP7	Yes
r17.5	r16.5 SP7	Yes
r17	r16.5	Yes
r16.5	r16.5, r16, r16 SP1, r16 SP2, r15, r15.1, r15.2	Yes

In addition, decide whether you will install the Arcserve Backup Client Agent on the Master server.

The benefits of installing the Client Agent on the Master server are as follows:

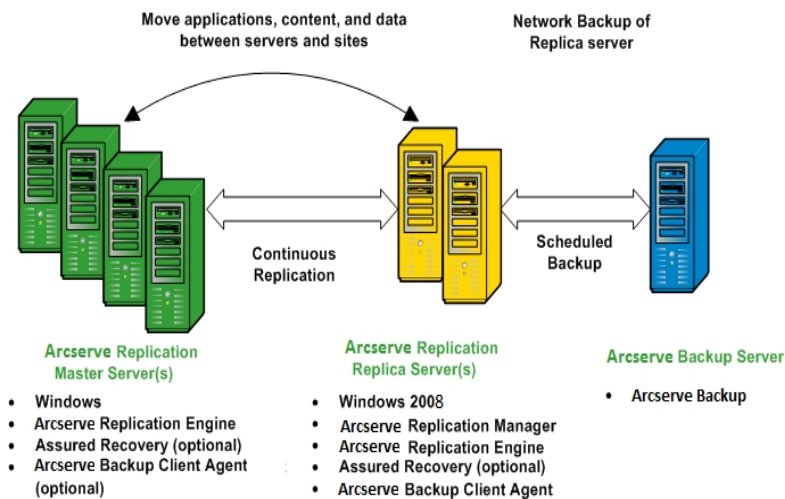
- Direct restores from the Arcserve Backup server to the Master server.
If the Client Agent is not installed on the Master server, you have to restore the Replica server and then perform an Arcserve Replication "reverse replication" (from Replica to Master) to help ensure that the Master server is fully up-to-date when you switch back to it.
- Bare-metal recovery, which allows the Master server to be rebuilt from scratch, including complete recovery of the server, with its applications and data, after a catastrophic failure or disaster.

Configuration with a Stand-alone Arcserve Backup Server

This setup involves a configuration where the Arcserve Backup server is installed on a separate standalone machine from the Replica server. The benefits of this configuration are as follows:

- Backups do not affect the functionality of the Replica server.
- Replication is quicker because the Arcserve Backup processes are running on a separate machine.
- This configuration meets the requirements for performing centralized backups for multiple remote branch offices (RBO).

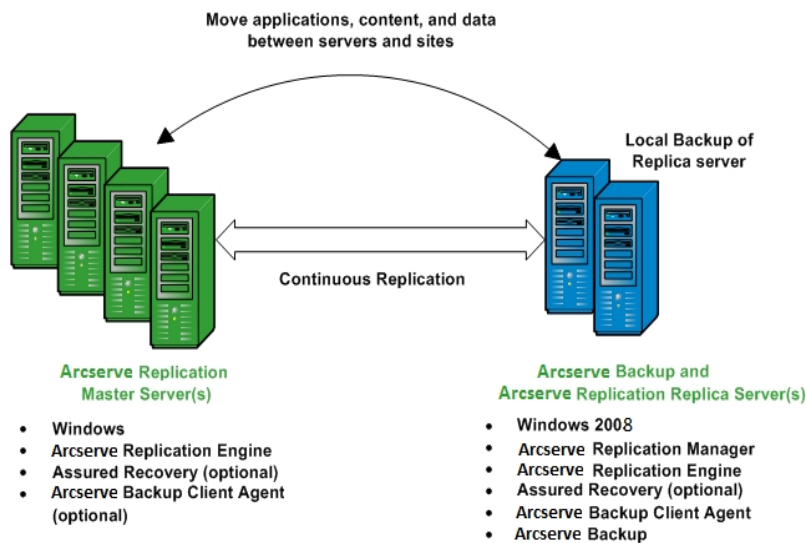
The following diagram shows the requirements of this configuration:



Configuration with an Arcserve Backup Server Installed on a Replica Server-Arcserve Replication

This setup involves a configuration where the Arcserve Backup server is installed on the same machine as the Replica server. This configuration lets you perform backups and restore jobs more quickly because these actions are performed local to the Replica server.

The following diagram shows the requirements of this configuration:



Remote Branch Office (RBO) Configurations

Arcserve Backup and Arcserve Replication together lets you replicate and back up data from remote branch office (RBO) servers to a central data center (CDC).

These RBO servers can be externally connected via through WAN or a LAN.


The advantages of RBOs using a CDC are as follows:

- Better disaster recovery and continuous application availability for your branch office servers
- Centralized and consolidated backups of multiple branch offices at a single facility
- Reduced need for competent and sophisticated IT support at every location
- Reduced cost associated with installing and maintaining multiple servers, storage devices, and applications in many different locations
- Reduced security risk associated with tape transport, and offsets tape media and handling costs.

Remote branch offices can be set up for replication and backup in a variety of configurations, depending upon your requirements and capabilities. The following diagrams provide a few examples of RBO configurations.

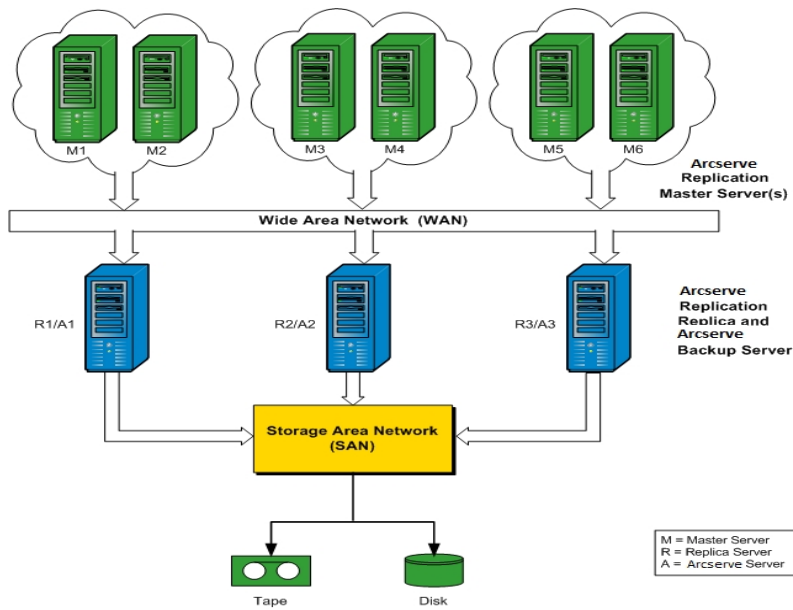
RBO Configuration Example 1

This example shows an RBO configuration of multiple Master servers being replicated to multiple Replica servers, and then backed up from a stand-alone Arcserve Backup server.




RBO Configuration Example 2

This example shows an RBO configuration of multiple Master servers being replicated to multiple Replica servers, and then being backed up from an Arcserve Backup server installed on each Replica server to a common library. In addition to the other advantages associated with remote branch offices using a CDC, this configuration also provides the advantage of local backups of all Arcserve Backup servers.



RBO Configuration Example 3

This example shows an RBO configuration of multiple Master servers being replicated to a single Replica server, and then being backed up from an Arcserve Backup server installed on the Replica server. The disadvantage of this configuration is the bottleneck condition caused by the Arcserve Replication Replica and Arcserve Backup server processing all replications and backups of multiple servers from multiple sites.



Chapter 3: Performing Integrated Backup Jobs

This section contains the following topics:

Integrated Backup Jobs	38
Create an Arcserve Replication Scenario	39
Run an Arcserve Replication Scenario	42
Create and Run a Backup Job	45

Integrated Backup Jobs

The integration between Arcserve Backup and Arcserve Replication provides the dual benefit of continuous data protection and backup of this protected data. Through this integration, Arcserve Replication provides real-time, continuous data replication from the Master server to the Replica server, while Arcserve Backup then backs up this replicated data from the Replica server to the Arcserve Backup server.

The process of performing an integrated backup involves the following operations:

- Creating a scenario
- Running a scenario
- Creating and Running a backup job

Create an Arcserve Replication Scenario

An Arcserve Replication scenario is the basis for managing the operation of the system. A scenario is a structure that describes what applications and data are to be protected, where they are located, where the data is to be replicated, and other scenario-specific options. Before performing a backup of an Arcserve Replication scenario, you must create a scenario to be backed up.

To create an Arcserve Replication scenario

1. Launch the Arcserve Replication Manager from either the Arcserve Replication start menu or from the Arcserve Backup Quick Start menu.

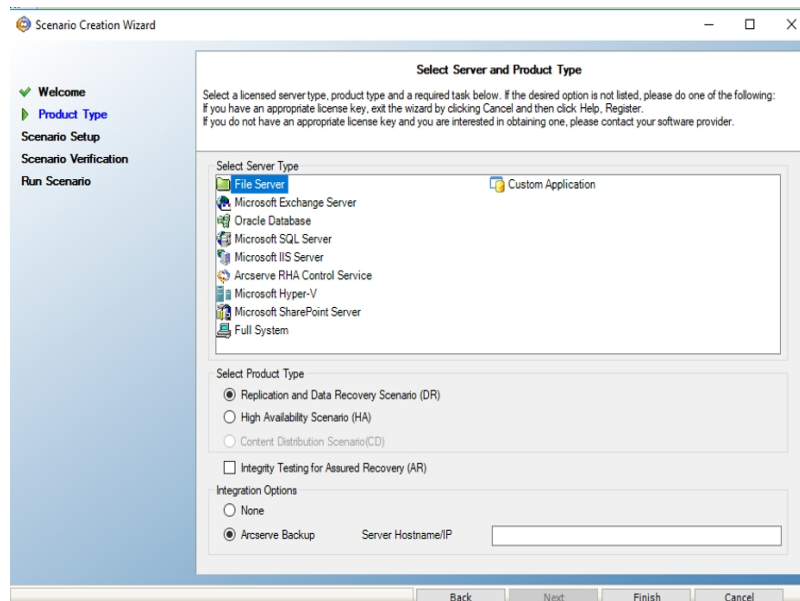
Note: For more information about replication scenarios and scenario creation options, see the [Arcserve RHA Administration Guide](#).

2. Click the New icon on the toolbar.

The Welcome to the New Scenario Wizard screen appears.

3. Select the Create New Scenario option and click **Next** to create a new scenario.

The Select Server and Product Type screen appears.

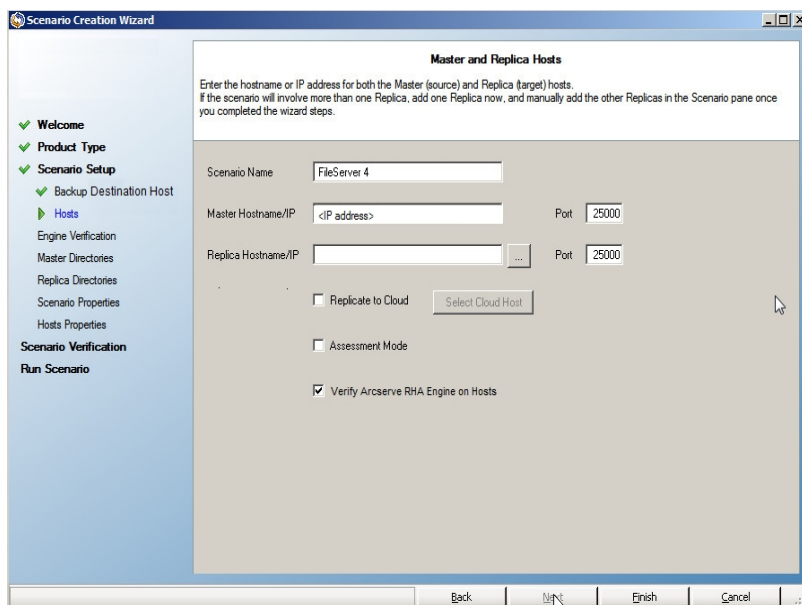


4. Select the Arcserve Backup option, select whether or not to include Integrity Testing, and enter the name of the Arcserve Backup server where the scenario will be backed up to. Select the appropriate Server Type, Product Type, and Assured Recovery (if available).

Note: Arcserve Backup supports File Server, Microsoft Exchange, and SQL Server scenarios only.

5. Click **Next**.

The Master and Replica Hosts screen appears.



6. Enter the Scenario name for the scenario being created and the hostname or IP address for both the Master server and Replica server. You can also use the browse buttons next to each hostname field to select the corresponding locations.

The scenario to be backed up has been created and the related information has been inserted into the Arcserve Backup database.

Note: For an existing scenario, you can update the related information in the Arcserve Backup database by using the Update Arcserve Backup Server option, which is accessible from the Tools menu of the Arcserve Replication Manager.

7. Click Next.

Note: If File Server or Microsoft Exchange Server was selected as the scenario type, continue with the same scenario creation process as any other Arcserve Replication scenario. For additional create scenario procedures, see the [Arcserve RHA Administration Guide](#).

If you selected SQL Server as the scenario type, the Master Configuration screen appears.

The Master Configuration screen displays all SQL Server databases for the master host, along with a check box for the "Replicate new user created databases in listed root directories" option.

With this option checked, if a new database is created in the SQL Server root directory after the scenario has been created, Arcserve Replication will automatically begin to replicate the new database to the Replica server. However, because the

newly created database has not been included in the Arcserve Backup database, the new database will not be backed up by Arcserve Backup. To allow the new database to be backed up, you must modify the scenario by running the Arcserve Replication auto-discovery function so that the new database will be recognized and included in the Arcserve Backup database.

Note: The Arcserve Replication auto-discovery function automatically discovers all database objects, related files, and directories on your database or mail server (local or on a network). For additional procedures about the auto discovery function, see the [Arcserve RHA Administration Guide](#).

8. After the scenario is properly configured, click Next and continue with the same scenario creation process as any other Arcserve Replication scenario.

Note: For additional create scenario procedures, see the [Arcserve RHA Administration Guide](#).

Important! If you make the following changes to your Master server after the scenario has been added to Arcserve Backup, you need to perform an additional procedure to make sure the changes are recognized by Arcserve Backup.

- Add a new database to a Microsoft SQL Server Master server
- Add a new storage group to a Microsoft Exchange Server Master server

After you change a Master server

1. Stop the running scenario.
2. Select Auto-Discovery and save.
3. Restart the scenario.


The new database or storage group should be now displayed in the Arcserve Backup Manager.

Run an Arcserve Replication Scenario

Arcserve Replication creates and maintains backups in the context of user-defined scenarios. Before you back up an Arcserve Replication scenario, you must run the scenario so that it is added to the Arcserve Backup database.

To run an Arcserve Replication scenario

- From the Arcserve Replication Manager interface, select the scenario to be backed up.
 - The scenarios and scenario status appear in the left pane.
 - The corresponding framework displaying directories (and subdirectories) and the files in those directories appear in the right pane.
 - The Events pane at the bottom displays information about significant events, warnings, and errors received from the host.




1 Scenario

2 Framework

3 Events

- Click the Run scenario icon on the toolbar.

The Arcserve Backup Server Connection dialog is displayed with the name of the server.




3. On the Arcserve Backup Server Connection dialog, enter the proper username and password to allow the Arcserve Backup server to communicate with the agent on the Replica server.

The selected scenario is now added to the Arcserve Backup database and can be backed up.

Note: You can perform backups only when Arcserve Replication is in the replication mode (green arrow icon next to scenario name).

The Arcserve Backup Activity Log is updated to reflect any Arcserve Replication operations related to creating, deleting, or modifying scenario records in the Arcserve Backup database.



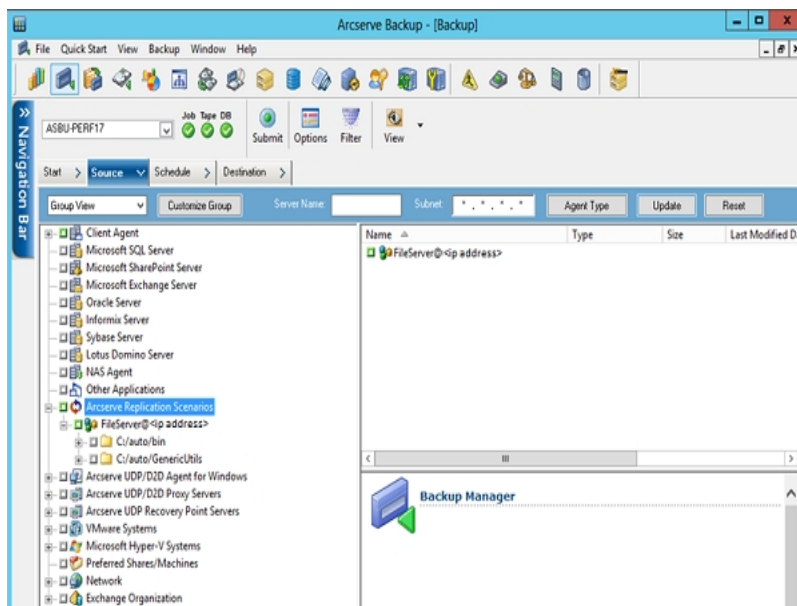
Create and Run a Backup Job

After an Arcserve Replication scenario has been added to the Arcserve Backup database a backup job can be created.

To create and run a backup job

1. From Arcserve Backup, access the Backup Manager and select the Source tab.

The left pane of the Source tab lists all of the Arcserve Replication scenarios that have been registered in the Arcserve Backup database and are candidates for being backed up.



2. Expand the Arcserve Replication Scenarios directory and select the scenario to be backed up.

The Security dialog appears, prompting you to provide the user name and password to log in into the Arcserve Backup Client Agent running on the Replica server.

3. Enter the user name and password and click OK.

Note: Arcserve Backup does not support logging in to systems with passwords that are greater than 23 characters. If the password on the system you are attempting to log in to is greater than 23 characters, you must modify the password on the agent system such that it is 23 characters or less, and then you can log in to the agent system.

If the credentials are accepted, you can select a scenario for backup.

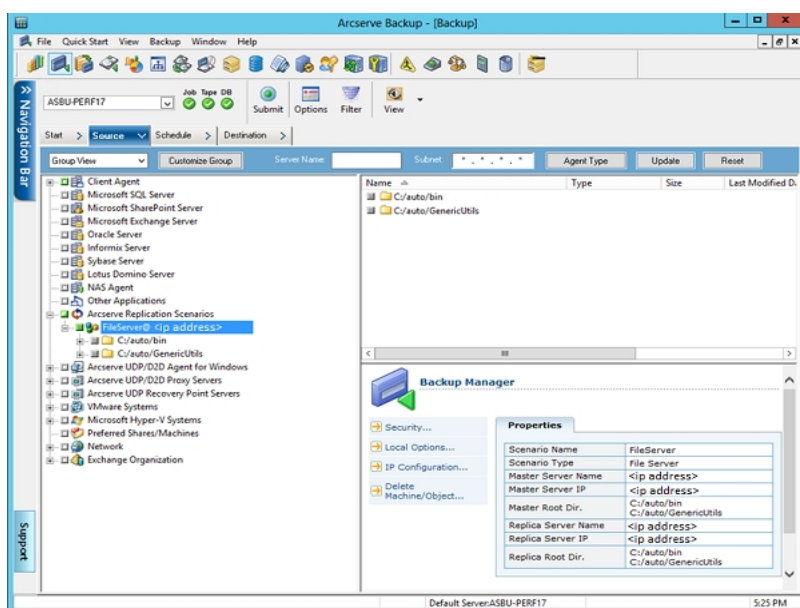
4. Expand the scenario to view the entities contained within the scenario to be backed up.

The scenario is expanded to display the associated entities. You can only view individual entities and not select them for backup. Supported scenarios include File Server, MS Exchange, and SQL Server.

Note: Depending upon the type of scenario stored in the Arcserve Backup server (File Server, MS Exchange, or SQL Server), different entities within each scenario are displayed.

5. Select the scenario you want to back up.

The marker box next to the selected scenario is filled solid and the corresponding scenario properties are displayed in the right pane of the Backup Manager. All backups are full backups (not incremental or differential).



The backup job creates a session for each entity that constitutes a scenario. The level of granularity for an entity depends on the type of scenario (File Server, MS Exchange, or SQL Server).

For multistreaming backup jobs, each Arcserve Replication scenario is backed up as a child job. If one node contains multiple scenarios, the master job will split them so that each child job will back up one scenario.

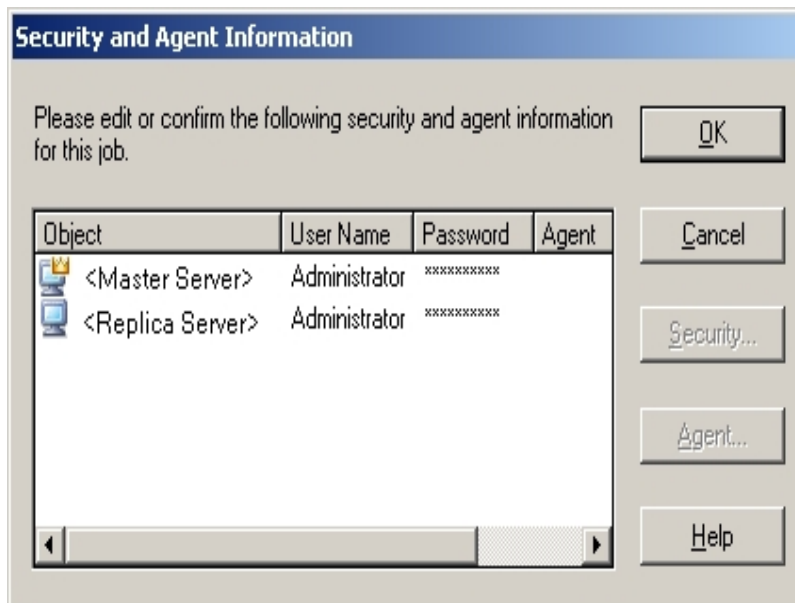
Note: You can back up entire scenarios only.



6. Make the relevant selections for Staging, Destination, Schedule, Global Options, and any other backup-related options. You can select multiple Arcserve Replication scenarios or include non-Arcserve Replication scenarios in the backup.


Note: For more information about these backup options, see the [Arcserve RHA Administration Guide](#).

7. From the Backup Manager, click Submit to initiate the backup.

The Security and Agent Information dialog appears and displays information about the selected scenario.



Each Arcserve Replication scenario displays two sets of credentials; one for the Master server and one for the associated Replica server. A graphic symbol next to the server name indicates if the server is a Master (active) server  or a Replica (standby) server .

 The Master server credentials are used to log in to the Arcserve Replication engine on the Master server, while the Replica server credentials are used to log in to the Arcserve Backup Client Agent running on the Replica server.

Note: The Master server credentials cannot be verified by the Preflight Checklist (PFC) utility.

8. Select the applicable server and click Security.

The Security dialog will appear for the selected server.

9. Enter the user name and password, and click OK.

Note: Arcserve Backup does not support logging in to systems with passwords that are greater than 23 characters. If the password on the system you are attempting to log in to is greater than 23 characters, you must modify the password on the agent system such that it is 23 characters or less, and then you can log in to the agent system.

The Submit Job dialog appears.

Continue the backup procedure as usual. For more information about this procedure, see the [Arcserve Backup Administration Guide](#).

10. When the backup procedure is completed, click OK to submit the backup job

Depending upon selected options, the backup job is either run immediately or entered in the job queue to be executed at the scheduled time.

Important! If you access the Arcserve Replication Manager while the backup job is running, a pop-up message indicates that the specified scenario is locked from another host and asks you to press OK to take control. If you press OK, the backup job will fail because Arcserve Backup is no longer controlling the scenario. To successfully complete the backup job, you must not open the Arcserve Replication Manager, click Cancel, or ignore the message and not press OK.

Note: Arcserve Backup cannot back up servers participating in Arcserve High Availability (HA) scenarios if a switchover has occurred. Backup jobs will fail while a scenario is running in the backward direction. After the switchback occurs and the scenario is running in the forward direction, Arcserve Backup can then successfully back up these servers.

Chapter 4: Performing Integrated Restore Jobs

This section contains the following topics:

Integrated Restore Jobs	50
Restore Using Arcserve Replication Failover	51
Restore Using Arcserve Replication Data Rewind	52
Restore Using Arcserve Backup	53

Integrated Restore Jobs

The integration between Arcserve Backup and Arcserve Replication lets you restore backed-up data from a variety of sources with a variety of methods.

The process of performing an integrated restore involves the following operations:

- Restore using Arcserve Replication Failover
- Restore using Arcserve Replication Data Rewind
- Restore using Arcserve Backup

More information:

[Restore Using Arcserve Replication Failover](#)

[Restore Using Arcserve Replication Data Rewind](#)

Restore Using Arcserve Replication Failover

Failover is a restore method that detects when the protected application on the Master server fails and switches (either automatically or manually activated) to a designated Replica server with essentially zero loss of data and time. If a Master server fails or must be shut down for maintenance, a synchronized replica (locally or in a remote location) will instantly and automatically take its place.

The advantage of an Arcserve Replication restore using the failover feature is that data processing resumes immediately if there is a hardware failure on the Master server. You can recover data almost immediately from the Replica server, without any disruptions or loss of data or service.

Automatic failover is designed to allow applications that are running on the Master server to automatically switch to the Replica server. This process can either be completely transparent or it can be configured to require user intervention.

Manual failover can be initiated for a number of reasons (usually for maintenance purposes), but it still results in the switching of the application processes from the Master server to the Replica server.

Note: For detailed procedures about recovering lost data from a Replica server using the failover feature, see the [Arcserve RHA Administration Guide](#).

Restore Using Arcserve Replication Data Rewind

Data Rewind is a restore method lets you rewind files to a point in time (rewind point) before they were corrupted. Because replication continuously updates source data to the Replica server, the Replica server always holds the same data as in the Master server. In the case of data corruption, recovering the latest files from the Replica server will not help, because that data in the Replica server is also corrupted. These rewind points serve as checkpoints in the Rewind log that mark an event or operation. The actual information stored includes the operation that will undo the event if the rewind point is activated. Data recovery uses these rewind points or bookmarks in the Rewind log to reset the current data back to a previous state. Because this built-in rewind technology occurs on the Replica server, any "data rewinding" to a previous state can only be performed on the Replica server.

The advantages of an Arcserve Replication restore using the data rewind feature are very quick recovery, extreme granularity of the recoverable data, and application-aware replication and recovery.

Note: For detailed procedures about recovering lost data from a Replica server using the Data Rewind feature, see the [Arcserve RHA Administration Guide](#).

Restore Using Arcserve Backup

You can recover data that was replicated from the Arcserve Replication Master server to the Arcserve Replication Replica server and then backed up with Arcserve Backup using the following methods:

- Restore by Session
- Restore by Tree
- Restore by Query

When restoring to a Replica server, ensure that the corresponding application is not running to avoid attempting to restore files that are currently in use. When restoring to the Master server, ensure that the Arcserve Backup Client Agent is installed and running.

Note: For SQL Server and MS Exchange scenarios, only full session restores are supported. For File Server scenarios, more granular restores to the specific files, directories, or volumes are supported.

In addition, regardless of the restore method being used, you can also make the relevant selections for Schedule, Global Options, and any other restore-related options. For more information about these restore options, see the [Arcserve Backup Administration Guide](#).

More information:

[Set Arcserve Replication-specific Global Restore Option](#)

[Restore by Session](#)

[Restore by Tree](#)

[Restore by Query](#)

Set Arcserve Replication-specific Global Restore Option

The Restore Manager contains a Arcserve Replication-specific global restore option that can be set regardless of which restore method is selected.

To set the Arcserve Replication-specific global restore option

1. From the Restore Manager window, click the Options toolbar button.


The Global Options dialog appears.

2. Select the Operation tab.

The Operation dialog appears, displaying the Arcserve Replication-specific "Continue the restore job even when the scenario cannot be stopped" option.

By default, this option is not checked, indicating that if Arcserve Replication cannot stop the scenario during the restore process, the job will fail. If you check this option, Arcserve Replication will still try to stop the scenario; however, if the scenario cannot be stopped, the restore job will continue. This option has the following advantage and disadvantage:

- The advantage of checking this option is that you are more likely to have a successful restore.
- The disadvantage of checking this option is that with the scenario running and continuous replication being performed, any problems that are contained on the Master server will overwrite the restored data on the Replica server.



Restore by Session

The Restore by Session method lets you select the session and the files and directories you want to restore. Use this method when you know the media name, but are not certain about the session you want to restore. This view uses the Arcserve Backup database; if the database engine is stopped, this method of restore will not work.


To restore a backup job by session

1. From Arcserve Backup, access the Restore Manager, select the Source tab, and choose Restore by Session from the Source View drop-down list.

The left pane of the Restore Manager lists all of the Arcserve Replication Scenarios that have been backed up and are candidates for being restored.

2. Select the session to be restored.

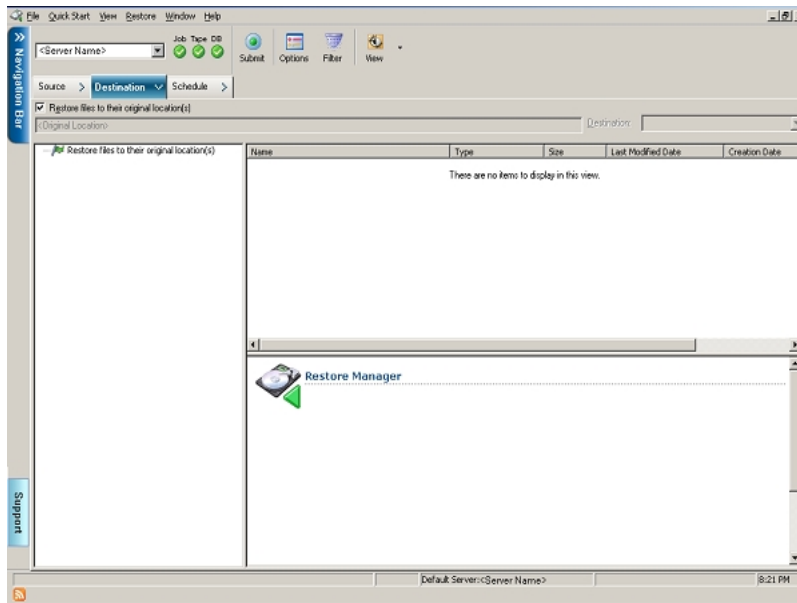
The corresponding session content and properties are displayed in the right panes of the Restore Manager.



3. Select the Destination tab and specify the destination where you want the restored files to go.

You can restore source files to the same directory structure that they originated from (original location) or to any other location that you specify (alternate loc-


ation).



4. To restore the files to their original location (default option), perform the following tasks:

- a. Ensure the Restore files to their original location(s) option check box is checked and click the Submit button to run the job.

The Session User Name and Password dialog appears.



For SQL Server and MS Exchange, if the corresponding application is running and you restore an Arcserve Replication scenario to the original location, the restored files are temporarily created with a .TMP extension. After the

restore job is successful, you are prompted to reboot the server to overwrite and replace the existing active files. When the server is rebooted, the restored files are merged into the original database. If the corresponding application is not running, temporary files will not be created during the restore to original location process and you do not have to reboot the server.

Important! Rebooting a SQL Server installed on a Cluster environment is not feasible, so you must stop the application resource before performing a restore to the original location.

- b. Select the server (Master or Replica) to restore the files to, and either double-click the selected row or click Edit.

The Enter User Name and Password for Arcserve Replication Session dialog appears.

- c. Select the Restore Option as either Restore to Master Server or Restore to Replica Server.

Restoring to a Replica server is usually more efficient and does not interrupt operation of the production (Master) server. Select Restore to a Master server only if you need a faster recovery or if you need to reconstruct the Master server environment (if corrupted).

For all releases of Arcserve Backup, the Replica server is selected by default.

Note: When restoring to a Replica server, ensure that the corresponding application is not running to avoid attempting to restore files that are currently in use. When restoring to the Master server, ensure that the Arcserve Backup Client Agent is installed and running.

- d. Enter the server security credentials (User Name and Password) and Session Password.

Note: Password Management provides the option to encrypt session passwords during backup and eliminates the need to repeatedly provide passwords. During a backup job submission, the passwords are stored in encrypted form and will be automatically used during restore. For more information about Password management, see the [Arcserve Backup Administration Guide](#).

- ◆ If you chose to restore to a Replica server, you must provide the security credentials to access both the Master server and the Replica server.
- ◆ If you chose to restore to a Master server, you must provide only the security credentials to access the Master server.

- e. Click OK.

The Submit Job dialog appears.

- f. Continue with the same restore procedure as detailed for any other Arcserve Backup restore job. For more information about this procedure, see the [Arcserve Backup Administration Guide](#).

5. To restore the files to an alternate location, perform the following tasks:


- a. Remove the check from the Restore files to their original location(s) option check box and select a destination folder.

The destination where you want the restored files to go is now specified.

Important! An Arcserve Replication restore is supported only through a Windows Systems Client Agent. As a result, the alternate restore location must be a destination machine or volume that is under the Windows Systems Client Agent tree. If you try to restore to the Server tree or any other tree not under the Client Agent tree, the restore job will fail. If necessary you can add a new Client Agent machine. For procedures on adding a new client object, see the [Administration Guide](#).

- b. Click Submit to run the job.

The Session User Name and Password dialog appears.



- c. Select the machine to restore the files to, and either double-click the selected row or click the Edit.

The Enter User Name and Password dialog appears.

- d. Enter the security credentials (User Name and Password), and click OK.

The Submit Job dialog appears.

- e. Continue with the same restore procedure as detailed for any other Arcserve Backup restore job. For more information about this procedure, see the [Arcserve Backup Administration Guide](#).

Restore by Tree

The Restore by Tree method restores a specific directory or drive from a display of files and directories that were backed up with Arcserve Backup. Use this method when you do not know which media contains the data you need, but you know the machine from which the backup originated.

The Restore by Tree view displays only the last instance of a backup. To view and access all other instances, select the object that you want to restore and click the Recovery Point button. The Restore by Tree view displays only the Master server specific view.

To restore a backup job by tree

1. From Arcserve Backup, access the Restore Manager, select the Source tab, and choose Restore by Tree from the Source View drop-down list.

Note: If necessary, you can choose a Recovery Point date and select a different session to restore.

The left pane of the Restore Manager lists the volumes, drives, directories, and files that have been backed up and are candidates for being restored.

2. Select the data to be restored.

The corresponding content and properties are displayed in the right panes of the Restore Manager.

3. Select the Destination tab and specify the destination where you want the restored files to go. Source files can be restored to the same directory structure that they originated from (original location) or to any other location that you specify (alternate location).

The destination where you want the restored files to go is now specified.

The Session User Name and Password dialog appears.

For SQL Server and MS Exchange, if the corresponding application is running and you restore an Arcserve Replication scenario to the original location, the restored files are temporarily created with a .TMP extension. After the restore job is successful, you are prompted to reboot the server to overwrite and replace the existing active files. When the server is rebooted, the restored files are merged into the original database. If the corresponding application is not running, temporary files will not be created during the restore to original location process and you do not have to reboot the server.

An Arcserve Replication restore is only supported through a Windows Systems Client Agent. As a result, the alternate restore location must be a destination machine

or volume that is under the Windows Systems Client Agent tree. If you try to restore to the Server tree or any other tree not under the Client Agent tree, the restore job will fail. If necessary you can add a new Client Agent machine. For procedures on adding a new client object, see the [Administration Guide](#).

Important! Rebooting a SQL Server installed on a cluster environment is not feasible, so you must stop the application resource before performing a restore to the original location.

4. Continue with the same restore procedure as detailed for the Restore by Session method.

Restore by Query

The Restore by Query method restores files based on the search pattern used to locate the names of the files or directories. Use this method when you know the name of the file or directory you want to restore, but do not know the machine it was backed up from or the media it was backed up to. This view uses the Arcserve Backup database.

Note: The Restore by Query method only supports File Server scenarios.

To restore a backup job by query

1. From Arcserve Backup, access the Restore Manager, select the Source tab, and choose Restore by Query from the Source View drop-down menu.

The top pane of the Restore Manager displays fields let you enter the search criteria for the scenarios that have been backed up and are candidates for being restored.

2. Specify the search criteria and click Query.

The bottom pane of the Restore Manager displays all the returned items that match the query criteria.

3. Select the files or directories that you want to restore and click Submit to run the job.

If the selected file or directory is an Arcserve Replication scenario, the Enter User Name and Password for Arcserve Replication Session dialog appears.

Restore Option

Restore to Master Server <Master Server>

Restore to Replica Server <Replica Server>

OK

Cancel

Help

Master Server User Name

Master Server Password

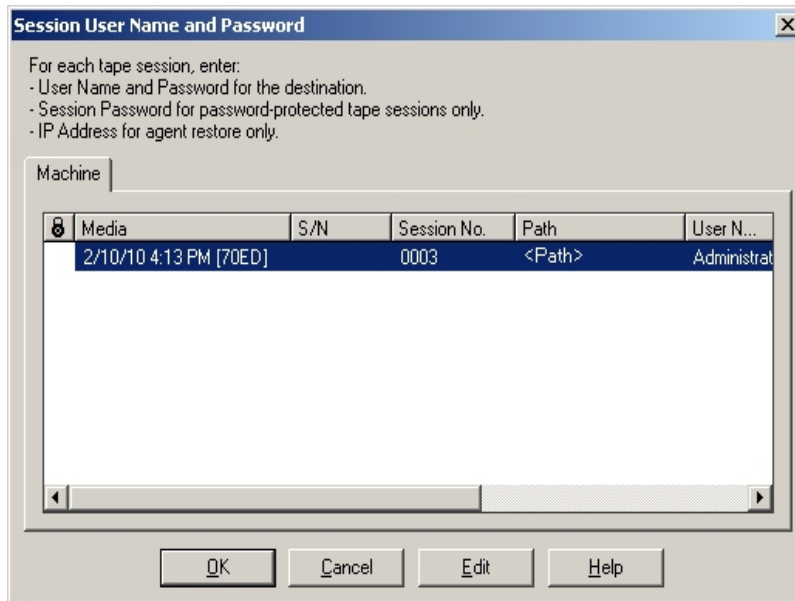
Replica Server User Name

Replica Server Password

Session Password

Apply [Option, User Name and Password] to all rows.

If the selected file or directory is not an Arcserve Replication scenario, the Session User Name and Password dialog appears.



4. Continue with the same restore procedure as detailed for the Restore by Session method.

Restore MS Exchange Scenario to Master Machine

When restoring backup data from a MS Exchange scenario to the Master server, the mailbox store is not automatically dismounted before the restore and mounted after the restore job is finished. You will need to perform the following procedure to restore an Exchange mailbox database to the Master server.

Restore an MS Exchange scenario to the Master server

1. From the Exchange Management Console, individually dismount each database in the mailbox database that you are restoring.

All databases in the mailbox database are dismounted.

2. Access the Microsoft Exchange Server Mailbox folder(s) and delete all mailbox database files from the folder.

All mailbox database files are deleted.

Note: If enough free disk space is available on your server, you can rename the folder instead of deleting the contents and then delete it after the restore job is successful.

3. Perform the restore to the Master server (using the Restore by Session, Restore by Tree, or Restore by Query method).
4. When the restore is successfully completed, return to the Exchange Management Console and mount each database in the mailbox database that was just restored.

Restore Microsoft SQL Server Database to Master Machine

When restoring a Microsoft SQL Server database to the Master server with Arcserve Backup, you may receive a "Windows could not start the SQL Server" error if running SQL Server 2008. This error is caused by the Network Service or Local Service accounts, which do not have access rights to the SQL log file on the Master.

There are three ways to avoid this issue:

- In Arcserve RHA, set the scenario Replicate ACL option to On. Doing so guarantees that all ACL information for the Master SQL data directories are retained after recovery, allowing the SQL service to operate.
- In Arcserve Backup, manually add the following user account after recovery to the ACL of the SQL data directories on the Master: `SQLServer-MSSQLUser$Computer Name$Instance Name`
- In Arcserve Backup disable the following restore option to NOT overwrite the ACL of the SQL data folders on the Master after restore: Arcserve Backup Manager, Global Options, Operation. Click to enable or disable the option, Restore and Preserve File Attributes and Security Information. Checking this option overwrites the ACL. Clearing the option does NOT overwrite the ACL.

Chapter 5: Monitoring Backup and Replication Jobs

This section contains the following topics:

Integrated Job Monitoring	68
Monitor Job Status Using Arcserve Backup	69
Monitor Job Status Using Arcserve Replication	72
Alert Notification	75
Report Generation	78

Integrated Job Monitoring

Monitoring the status of backup and replication jobs consists of a combination of real-time event monitoring, generated alerts, and various reports. You can monitor an entire integrated backup process using Arcserve Backup and Arcserve Replication.

Monitor Job Status Using Arcserve Backup

The procedure for monitoring the backup process does not change for the integrated environment. For more information about monitoring the backup process, see the [Arcserve Backup Administration Guide](#).

You monitor the integrated backup process from Arcserve Backup through the Job Status Manager. The Job Status Manager is a graphical tool that helps you to centrally manage Arcserve Backup servers across the enterprise and monitors all pending, completed, and active jobs from the Job Status Manager window. The Job Status Manager window contains a Job Queue tab and an Activity Log tab to monitor.

Job Queue Monitoring

The Job Queue tab on the right panel displays information about all jobs. Every time you run or schedule a job with the Arcserve Backup Backup Manager, you submit it to the Job Queue. Arcserve Backup continuously scans the Job Queue for jobs that are waiting to execute.

When a job is in the Arcserve Backup queue, it is listed with a status. The status can be one of the following:

Done 

Indicates a job that has already been executed with no repeat interval.

Ready 

Indicates a new one-time or repeating job waiting to be executed. For example, a backup job that runs every Friday.

Active 

Indicates a job that is currently being executed.

Hold 

Indicates a job that is in the queue waiting to be executed later.

When an integrated backup job is submitted, the scenario name is displayed in the Job Queue window, letting you monitor the progress of the job.

Activity Log Monitoring

The Activity Log tab on the right panel displays comprehensive information about all the operations performed by Arcserve Backup. The Activity Log provides an audit trail of every job that is run. For each job, the log includes the following information:


- Time the job started and ended
- Type of job
- Average throughput of the data
- Number of directories and files processed (backed up, restored, or copied)
- Job session number and job ID
- Result of the job
- Errors and warnings that occurred

When an integrated backup job is submitted, the scenario name is displayed in the Activity Log window, letting you monitor the progress of the job.

Monitor Job Status Using Arcserve Replication

The procedure for monitoring the replication process does not change for the integrated environment. For more information about monitoring the replication process, see the [Arcserve RHA Administration Guide](#).

You can monitor the replication process from the Arcserve Replication Manager after a scenario is running. Monitoring lets you view state information, statistics, and events. The Arcserve Replication Manager main window is comprised of three subordinate panes; Scenario, Framework, and Events.



1 Scenario 2 Framework 3 Events

The Scenario pane displays each host and the corresponding status of the replication process. The replication status is reflected by one of the following icons to the left of the scenario name:



Running

Indicates that the scenario replication process is running correctly.



Stopped

Indicates that the scenario has been created, but the replication process has been stopped or suspended.



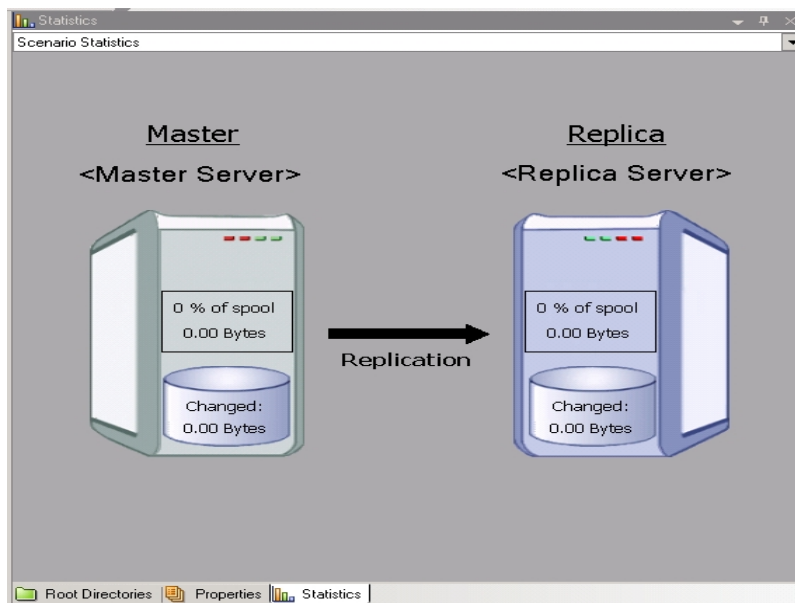
Problem

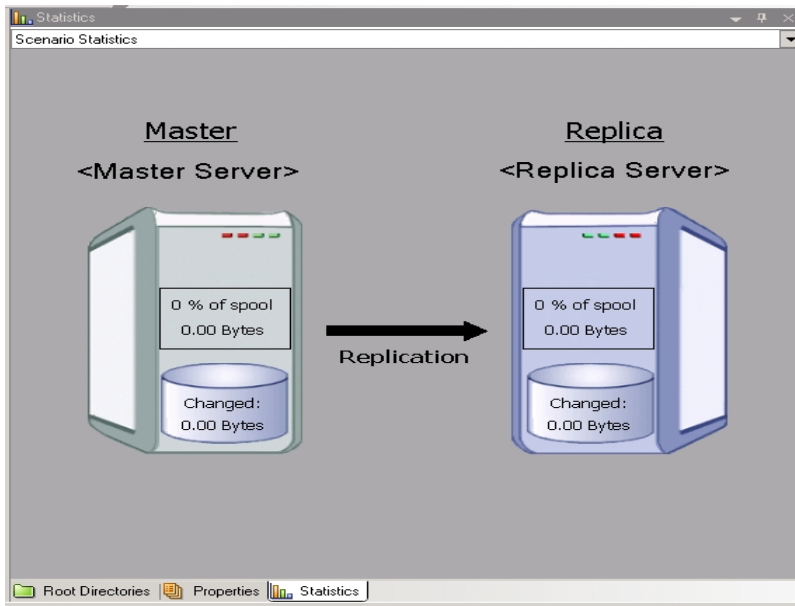
Indicates that there is a problem with the scenario. You can click on the scenario name to display any related error messages in the Events window.

Not Authorized

Indicates that incorrect or missing User Name or Password was given for the Master server.

The Framework pane displays the directories, subdirectories, and the files in those directories. The Framework pane displays two or three tabs depending upon the job status; Statistics, Directories, and Properties. The Statistics tab is only available when the replication process is running and provides information about the total amount of data for each root directory, recorded data for each Replica, and synchronization information. From a drop-down menu, you can select to view the Scenario Statistics (graphic overview of scenario status), the Master Statistics (table overview of Master Server status), or Replica Statistics (table overview of the Replica Server status).





The Events pane displays messages and general information (for example, that a directory is synchronized, a server is connected, synchronization has started or finished, and so on). This information is received from the servers participating in the running replication scenario. The information in the Events pane includes the server name and time, and a brief explanation of the event. Important events or error messages are displayed in bold letters. In addition, the Events pane also displays Arcserve Backup status for backup jobs that are initiated through Arcserve Replication.

The screenshot shows the 'Scenarios' pane with a table of running scenarios and the 'Events' pane with a log of events.

Scenario	State	Product	Server	Mode
SQL	Running	DR/AR/A...	SQL	Online
FileServer	Running	DR/AR/A...	FileServer	Online

ID	Sequence	Severity	Host/Scenario	Time	Event
SR00104	95	Significant	<Host Name>	2/10/2010 4:59:37 PM	Replication to replica <Server Name> resumed
IM00405	94	Info	FileServer	2/10/2010 4:59:36 PM	Posting Assured Recovery report created at '2/10/2010 16:59:36' to Reports
SR00392	92	Significant	<Host Name>	2/10/2010 4:59:36 PM	FileServer Integrity Testing on replica <Server Name> is finished
IR00199	91	Info	<Host Name>	2/10/2010 4:59:36 PM	Shadow Copy Id for volume C:\ is {4b18e04-bd00-46d3-84bd-61a295f999f3}
IR00197	90	Info	<Host Name>	2/10/2010 4:59:36 PM	Shadow Copy is built successfully
IR00175	89	Info	<Host Name>	2/10/2010 4:59:30 PM	Building Shadow Copy
IR00343	88	Info	<Host Name>	2/10/2010 4:59:30 PM	Replica <Server Name> suspended for integrity testing
IM00405	74	Info	FileServer	2/10/2010 4:21:24 PM	Posting Assured Recovery report created at '2/10/2010 16:21:24' to Reports
SR00104	72	Significant	<Host Name>	2/10/2010 4:21:25 PM	Replication to replica <Server Name> resumed
SR00392	71	Significant	<Host Name>	2/10/2010 4:21:24 PM	FileServer Integrity Testing on replica <Server Name> is finished
IR00199	70	Info	<Host Name>	2/10/2010 4:21:24 PM	Shadow Copy Id for volume C:\ is {4a7f659f-25f1-4483-bd3d-d93c-d28472e}
IR00197	69	Info	<Host Name>	2/10/2010 4:21:24 PM	Shadow Copy is built successfully
IR00175	68	Info	<Host Name>	2/10/2010 4:21:15 PM	Building Shadow Copy
IR00343	67	Info	<Host Name>	2/10/2010 4:21:15 PM	Replica <Server Name> suspended for integrity testing
SR00104	53	Significant	<Host Name>	2/10/2010 4:14:47 PM	Replication to replica <Server Name> resumed
IM00405	52	Info	FileServer	2/10/2010 4:14:46 PM	Posting Assured Recovery report created at '2/10/2010 16:14:47' to Reports
SR00392	50	Significant	<Host Name>	2/10/2010 4:14:46 PM	FileServer Integrity Testing on replica <Server Name> is finished

Alert Notification

The procedure for generating and receiving alerts does not change for the integrated environment.

- For more information about generated alerts during the backup process, see the [Arcserve Backup Administration Guide](#).
- For more information about generated alerts during the replication process, see the [Arcserve RHA Administration Guide](#).

Arcserve Backup Alerts

You can use the Alert notification system to send messages about events that appear in the Activity Log during your backup operation. In addition, you can also specify the method for receiving these alert notifications. You can choose one or more of the following events for which you want to be notified:

Job Completed Successfully

Indicates that all of the nodes and drives/shares were processed.

Job Incomplete

Indicates that some nodes, drives, shares, or files were missed.

Job Canceled by User

Indicates that the user canceled the job.

Job Failed

Indicates that the job was started but could not be completed.

Virus Detected

Indicates that a virus was detected in one of the files to be backed up.

Customized Event

Indicates that a customized event occurred. To specify this type of event, enter an error, warning, or notification code in the space below the Event drop-down list.

Arcserve Replication Alerts

All events are reported in real time to the Arcserve Replication Manager and can be integrated into the OS event logging system. The alerts can be automatically sent by email to a configured address and can also activate a notification script. Detailed real-time statistics are provided during synchronization and replication. When the event notification is configured for a scenario, the following conditions can trigger a notification:

Lost Connection

Indicates that the TCP connection does not function, or a network or server went down.

Queue overflow

Indicates that the quantity of data in the queue spool directory has exceeded its threshold value.

Other

Indicates any other error

Significant info

Indicates important information such as when synchronization is completed.

Report Generation

The procedure for generating reports does not change for the integrated environment.

- For more information about reports during the backup process, see the [Arcserve Backup Administration Guide](#).
- For more information about reports during the replication process, see the [Arcserve RHA Administration Guide](#).

Arcserve Backup Reports

The reports generated by the Arcserve Backup Report Manager and Arcserve Replication complement each other and provide you with a variety of reports based on the activity stored in the Arcserve Backup database. You can preview a report, print to a printer or file, and schedule when to generate a report. Arcserve Backup provides several standard reports that display general backup and restore activity and also lets you to create custom reports to meet your specific needs. You can use a report filter to select the backup media you want to include in the report.

For example, you can configure Arcserve Backup to automatically create and email a report for any failed backup jobs, or you can create a customized report that is generated to specifically target backup jobs that are initiated by Assured Recovery.

Note: For more information about reports during the backup or restore process, see the [Arcserve Backup Administration Guide](#).

Arcserve Replication Reports

Arcserve Replication can generate reports on the replication and synchronization processes. These reports can be stored on the Master, sent for display by the Arcserve Replication Manager, sent by email to a specified address, or they can trigger script execution.

The Arcserve Replication-generated reports include the following:

Scenario Reports

The Scenario reports include statistics from synchronization tasks, including how much changed data was replicated. You can configure this report so that it is displayed at the end of every synchronization process.

This report helps you to verify that all processes are running properly, and helps in reviewing how much data is continually changing.

Difference Reports

A Difference report compares the difference between the Master server and the Replica server and is generated for each replica at the end of the replication process. When a replication is suspended, all changes are spooled on the Master server until the replication is resumed. During this suspension, the Difference report displays how much data has changed. The Difference report can be produced at any time.

This report helps you determine how much data changes for a specific event.

Replication Report

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

This helps you get an overall view of how much data is changing in the environment.

Synchronization Report

Following synchronization, Arcserve Replication creates and opens a report listing the files that have been transferred. This report includes the sum total of the removed and modified files, and the bytes transferred, listing all related file names, file paths, and sizes.

This report helps in monitoring and managing data change and data growth in the environment.

Note: For more information about reports during the replication process, see the [*Arcserve RHA Administration Guide*](#).

Chapter 6: Troubleshooting

This section contains the following topics:

Integrated Troubleshooting	84
Error and Warning Messages	85

Integrated Troubleshooting

When a problem is detected, Arcserve Backup generates a message to help you identify the problem and resolve the problem. These messages are contained in the Activity Log and can be viewed from the Job Status Activity Log tab. From the Activity Log, you can double-click an error or warning message to display the message details.

Error and Warning Messages

Generated messages are categorized as either Error Messages or Warning Message, depending upon the severity of the resulting consequences. An Error Message is more severe and usually indicates a functionality problem that must be fixed before the job can continue. A Warning Message indicates a less severe problem that should be noted; however, the job can continue.

The generated message is formatted with some or all of the following information:

Message

Indicates the Warning or Error identification number prefixed by either a W (warning) or E (error), along with a corresponding error code or brief explanation generated by Arcserve Backup Agent message numbers are prefixed by either an AW (agent warning) or AE (agent error).

Module

Indicates the system component or area that produced the message.

Reason

Describes what may have caused the message to be generated.

Action

Suggests a possible resolution to the problem or a course of action you can take.

Chapter 7: Glossary

This section contains the following topics:

assured recovery	87
continuous data protection	87
data rewind	87
failover	87
master server	87
replica server	88
synchronization	88

assured recovery

Assured Recovery lets you perform a real test of your disaster recovery server by actually running the application, including modifying data, without impacting your production environment or your previously replicated data.

continuous data protection

Continuous data protection (CDP) is the ability to recover data not just to certain isolated previous states captured, for example, in a daily or weekly backup or snapshot, but to recover the data back to any point in time.

data rewind

Data rewind is a recovery method that allows rewinding files to a point in time (rewind point) before they were corrupted.

failover

Failover is a feature that detects when the protected application on the master server fails and switches (either automatically or manually activated) to a designated replica server with essentially zero loss of data and time.

master server

The master server is the active or production server that lets you actively change (read and write) data.

replica server

The replica server is the passive server. This is the server from which data cannot be changed (read only) in any way except through changes replicated from the master server.

synchronization

Synchronization is the process of bringing the data on the Replica server in sync with the data on the Master server.