

ARCserve[®] 2000

Administrator Guide

Version 7.0



This documentation and related computer software program (hereinafter referred to as the "Documentation") is for the end user's informational purposes only and is subject to change or withdrawal by Computer Associates International, Inc. ("CA") at any time.

THIS DOCUMENTATION MAY NOT BE COPIED, TRANSFERRED, REPRODUCED, DISCLOSED OR DUPLICATED, IN WHOLE OR IN PART, WITHOUT THE PRIOR WRITTEN CONSENT OF CA. THIS DOCUMENTATION IS PROPRIETARY INFORMATION OF CA AND PROTECTED BY THE COPYRIGHT LAWS OF THE UNITED STATES AND INTERNATIONAL TREATIES.

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

THE USE OF ANY PRODUCT REFERENCED IN THIS DOCUMENTATION AND THIS DOCUMENTATION IS GOVERNED BY THE END USER'S APPLICABLE LICENSE AGREEMENT.

The manufacturer of this documentation is Computer Associates International, Inc.

Provided with "Restricted Rights" as set forth in 48 C.F.R. Section 12.212, 48 C.F.R. Sections 52.227-19(c)(1) and (2) or DFARS Section 252.227.7013(c)(1)(ii) or applicable successor provisions.

© 1990-2000 Computer Associates International, Inc., One Computer Associates Plaza, Islandia, New York 11749. All rights reserved.

All product names referenced herein belong to their respective companies.



Contents

Chapter 1: Introducing ARCserve

Features	1-1
Windows 2000 Features	1-3
Components	1-4
Functionality	1-5
Managers	1-7
Wizards	1-8
Utilities	1-9
Recover DB	1-10
Merge	1-11
Scan	1-11
Compare	1-12
Count	1-12
Purge	1-12
Copy	1-13
Preferences	1-13

Chapter 2: Backing Up Your Data

The Backup Wizard	2-1
The Backup Manager	2-2
Backup Options	2-3
Backup Media	2-3
Options for First Backup Media	2-4
Options for Additional Media	2-5
Secure Your Backup Data	2-5
Verification	2-6
Retry	2-7
Operation	2-8
Pre/Post	2-9
Log	2-11

Virus	2-12
Alert	2-13
Replication	2-14
Media Exporting	2-15
Windows 2000	2-16
Backup Agent for Open Files	2-17
Disaster Recovery	2-18

Chapter 3: Restoring Data

The Restore Wizard	3-1
The Restore Manager	3-1
Source	3-3
Destination	3-4
Schedule	3-5
Restore Options	3-6
Backup Media	3-6
Destination	3-7
Operation	3-9
Pre/Post	3-9
Log	3-11
Virus	3-12
Alert	3-13
System State Restore Options	3-14

Chapter 4: Managing Jobs

The Job Status Manager	4-1
Manage Servers and Groups	4-2
Job Queue	4-3
Activity Log	4-4
Job Detail	4-5
The Job Log	4-5
Job Status	4-6
Backup Job Failure	4-7

Chapter 5: Filters, Scheduling and Job Scripts

Filters	5-1
Types of Filters	5-2
File Pattern Filter	5-2
Directory Filter	5-2
File Attributes Filter	5-3
File Modified Filters	5-3
File Created Filters	5-3
File Accessed Filters	5-4
File Size Filters	5-4
NDS Context & Object Filters	5-4
Scheduling	5-5
Job Scripts	5-6
Saving Scripts	5-6
Using an Existing Script	5-7

Chapter 6: Managing Devices and Media

Device Manager Basics	6-1
Device Management Functions for Libraries	6-2
Formatting Media	6-3
Erasing Media	6-4
Retensioning a Tape	6-4
Compressing Data on Tape	6-4
Ejecting Media	6-5
Rebuild	6-5
Inventory/Offline Slots	6-5
Mount/Dismount	6-5
Import/Export	6-6
Clean Tape Heads	6-6
Device Groups	6-7
Configuring Device Groups	6-7
Removable Drive Support	6-8
Configuring Libraries	6-9

Chapter 7: Media Pools and Rotation Schemes

Media Pool Basics	7-1
Media Pool Manager	7-3
Rotation and Grandfather-Father-Son (GFS) Media Pools	7-5
Simple Rotation Scheme	7-6
GFS Rotation Schemes	7-6
Incremental backups	7-7
Differential backups	7-7
GFS Media Naming	7-7
GFS Media Pools	7-8
Selecting GFS Rotation Schemes	7-11
Sample GFS Rotation	7-12

Chapter 8: Using RAID with ARCserve

RAID Levels	8-1
RAID Fault Tolerance	8-2
Configuring RAID Levels	8-3
Device Configuration	8-4
Backing Up and Restoring with RAID	8-4
Maximizing Performance with RAID	8-4
Troubleshooting	8-6

Chapter 9: Managing the Database and Reporting

The Database Manager	9-1
Selecting a Database View	9-2
Using a Centralized Database	9-4
Adding Members to an ARCserve Central Database	9-4
Configuring the Job Information Stored in Your Central Database	9-5
Using a Microsoft SQL Server	9-5
Transferring Data from a Raima Database to a SQL Database	9-5
SQL Connections	9-6
Run a Database Consistency Check	9-6
The Report Manager	9-6
The Activity Log	9-7
The Tape Log	9-7
The Job Log	9-7
Viewing Reports	9-8

Chapter 10: Administering the ARCserve Server

ARCserve Engines	10-1
Server Admin Engine tabs	10-2
Engine Status	10-4
InoculateIT	10-4
Configuring the ARCserve Engines	10-5
Additional Server Admin Functions	10-11

Chapter 11: Using the Alert Manager

Components	11-2
Alert Features	11-2
The Alert Manager	11-2
Configuring Alert	11-2
Alert's Broadcast Option	11-3
Alert's Pager Option	11-3
Interpreting the Pager Message	11-3
SNMP Option	11-4
Trouble Ticket Option	11-4
E-mail Option	11-4
Unicenter TNG Option	11-4
Sample TNG Alert Scenarios	11-5
Event Log Option	11-5
Event Priority	11-5
Alert's Activity and Event Logs	11-6
Testing the Recipients	11-6

Appendix A: ARCserve Cluster Support

Requirements for Installing ARCserve on a Cluster	A-2
Software Requirements	A-2
Hardware Requirements	A-3
Installing ARCserve on a Cluster	A-3
Registering ARCserve with the Cluster	A-4
ARCserve Services Fail-Over (Active/Active & Active/Passive)	A-5
ARCserve Job Fail-Over (Active/Passive)	A-5
ARCserveIT Checkpoints	A-6

Introducing ARCserve

ARCserve is a comprehensive, enterprise-level backup and restore management solution for multi-platform environments. It can back up and restore data from all the machines on your network, including machines running Windows NT/2000, 9x, 3.1x, Macintosh, UNIX, OS/2 or NetWare servers using optional Client Agents. ARCserve also provides media utilities for maintenance.

In addition, it provides backup management from one management console. ARCserve can support small- and large-scale networks, it can run on one machine or many, which can span thousands of miles, different support organizations, and different platforms.

ARCserve provides a comprehensive set of product documentation, including Windows online documentation and online help, Adobe Acrobat PDF files, and readme files containing up-to-date supplemental information.

Features

ARCserve provides the following features:

- **Easy-to-use GUI interface**—ARCserve administrators can use a simple yet advanced graphical interface to perform administration and job management functions. Explorer-like tree view, property dialog, drag-and-drop, and right-mouse click integrate with Windows 2000 advanced interface designs. You can configure and save column list display and layout (width, position, ordering, and display order).
- **ODBC database support**—ARCserve supports ODBC (SQL) database formats to use as its main database system.
- **Virus Scanning**— Since InoculateIT, Computer Associates premier anti-virus protection solution is bundled with ARCserve, ARCserve can automatically detect viruses while performing backup, restore, copy, and count operations.

- **Centralized Management of ARCserve**— This allows you to manage multiple ARCserve host servers from a single ARCserve Manager. You can back up, copy, and restore any machine and other servers in your network from one console. You can group preferred ARCserve Servers and view job status, and monitor active jobs as well as view activity logs. From the ARCserve Manager, you can also administer the various ARCserve host database systems from a single console. Using the Manager's reporting system, you can customize reports.
- **CRC Verification support**— When you back up your data, you can verify the individual data packets with CRC verification. CRC prevents data corruption by generating a numeric value for each packet, then checking that value against the data on the media.
- **Media Rotation Scheme** — ARCserve can schedule repeating backup jobs that are custom-made for your environment. ARCserve uses an automatic rotation scheme, which can be a simple rotation or a Grandfather-Father-Son (GFS) rotation to make scheduling easier. Large backup jobs will automatically span multiple media until the job is complete.
- **Media pool management**— The Media Pool Manager allows you to create and maintain media pools for a logical grouping of media. This allows you to efficiently schedule the maintenance and recycling of your media. ARCserve creates a catalog file on your media to improve performance for merge jobs and database backup jobs.
- **Changer support**— ARCserve provides built-in changer support for one tape drive. Using a changer enables you to group a range of slots together, which offers advantages such as tape spanning.
- **Backup Agents**— ARCserve can support a wide variety of complex databases and GroupWare applications such as Oracle, Informix, SQL, SAP R/3, MS-Exchange, and Lotus Notes. These agents are sold separately.
- **Client Agent support**— You can back up remote machines using the ARCserve Client Agents (for NetWare, Windows 9x, NT/2000, 3.1x, UNIX, OS/2, and Macintosh). Client Agents efficiently back up and restore complete system information for all the listed platforms employing PUSH technology. These agents are sold separately.
- **RAID (Redundant Array of Inexpensive Devices)**— This is a set of drives that behave as a single, fault-tolerant, high-performance drive. Fault tolerance allows you to continue your backup or restore, even if the media goes bad or if one drive fails. Using this feature improves backup and restore performance by distributing data over multiple drives. This option is sold separately.

- **Virtual Library Feature**— This is ARCserve's unique add-on feature to the tape library support, which allows ARCserve to dynamically configure any library to act as multiple independent libraries split up by drive(s). It also supports RAID within this feature.
- **Centralized Database**— You can use this to restore data and use media and session detail for generating ad hoc reports for the Central Database members. The members, which are subordinate ARCserve machines on your network, update the Central ARCserve Database with media session detail and information from their own databases.
- **Advanced event-driven notification system**— The Alert notification system informs you of key events in your operations. ALERT notifications can be used for specific non-job related events, such as the tape or database engine starting, error messages, or other events viewed through the Activity log.
- **Device configuration** – ARCserve allows you to select one of the following types of backup devices to configure: Tape and Optical, Removable Device, and RAID (Redundant Array of Inexpensive Devices). You can submit jobs concurrently when using multiple storage devices.
- **Security**— ARCserve provides intelligent, end-to-end remote and local client-to-server data compression. Compressing data increases the effectiveness of devices and decreases network traffic. All ARCserve backup jobs can be encrypted with a session/encryption password to enhance network security.
- **Pre/Post commands based upon job status**— ARCserve contains conditional pre-execution, job execution, and post-execution based on previous results (exit codes). For instance, you can run a detailed report on a job if it fails.

Windows 2000 Features

ARCserve provides the following features for Windows 2000:

- **System State Backup**— Windows 2000 has many components defined as the System State that must be backed up and restored together as a single entity. The components that comprise System State are: Registry, Active Directory, SYSVOL, Cluster Database, Certificate Server Database, COM+ registration database, Removable Storage Manager database, and Single Instance Storage.
- **Cluster fail-over Capabilities**— Enables the ARCserve services to handle its services fail-over and to enable ARCserve jobs to handle fail-over to another cluster node in case of fail-over. In the case of an ARCserve job fail-over, the failed job will be restarted at the other cluster node, and the failed job will be resumed at the volume/disk where it failed before by leveraging new ARCserve checkpoint technology.

- **Support for NTFS file system changes** – ARCserve seamlessly supports the changes in the file system with regards to Object level security. Among the new data streams that ARCserve is aware of are: Reparse Points, Object ID's, Sparse Files, Encrypted Files, Volume Mount Points, Disk Quotas, and Directory Junctions
- **Reparse Points** – Any file or empty directory can have a reparse tag attached to it. ARCserve checks if the object to be backed up has a reparse tag and safely backs up and restores the reparse data. This ensures a proper restore of any directory and volume links that the might have defined.
- **Restoring Active Files** – ARCserve takes advantage of the facility that the Win2000 system exposes to safely backup and restore active files. Files that are found active during a restore are written to a temporary file and the restore is completed by requesting the system to move the file to its original location during the next reboot.
- **Object ID's** – Object ID's can be placed on some files. This means that these files can then be accessed using path or by referring to its object ID. ARCserve ensures this by safely backing up and restoring the new stream that defines the Object ID.
- **Sparse Files** – Sparse streams allow for de allocations to occur anywhere in the data stream. This means that a 100 GB file can contain only 1MB of allocated data. ARCserve supports backing up and restoring streams and attributes specific to such files.
- **Encrypted Files** – Windows 2000 allows file encryption at a file system level. ARCserve backs up and restores the raw data and of the encrypted file.
- **Volume Mount Points** – A mount point is a reparse point that can be placed on any directory. Placing a mount point on a directory maps the volume under a directory on another volume. ARCserve supports backing up and restoring such mount points. It implements some intelligence to detect if the volume mapped under the directory is already a part of the backup job and chooses to ignore or process the mapping.

Components

ARCserve consists of two main components that work together to back up, copy, and restore your data. The components are the:

- ARCserve Manager
- ARCserve Server Admin

ARCserve Manager

Use the ARCserve Manager to submit your backup and restore jobs, manage your database, and search reports. It can be installed on the same machine as the ARCserve Server or onto a different machine, i.e. a Windows 2000 workstation.

ARCserve Server
Admin

Control ARCserve services through the Server Admin. These core services are the Job, Tape, and Database Engines:

- The Job Engine processes your jobs at their designated date and time. It scans the job queue for a job that is ready to run, then sends it to the appropriate handler.
- The Tape Engine is responsible for communicating with and controlling your storage devices. The Tape Engine selects the device needed for an ARCserve job.
- By default, all operations within ARCserve are recorded by the Database Engine. The Database Engine is responsible for maintaining a history of:
 - Files, directories, drives, and machines that ARCserve has backed up or copied.
 - Information about jobs that have been processed by ARCserve, such as the job type, logs, the final result, and the start and end time.
 - Media used by ARCserve, such as its type, its name, the date it was first formatted, the date it expires, and the sessions on it.

Functionality

ARCserve furnishes the components, functions and utilities required by Network Managers to obtain and actively manage network backups. Start the ARCserve Manager by selecting the ARCserve Manager icon from its program group. When you launch ARCserve, the ARCserve Home Page appears. From the Home Page you can select options for all of ARCserve's functions.



There are three main sections on the ARCserve Home Page:

- Contents— Quickly gets you from page to page to access the features or functions of ARCserve.
- Most Recently Used — ARCserve allows you to go directly to the most recently used screens.
- Main Window — Provides links to the Quick Access, Wizards, News, and Support pages.

The following describes the pages available in the ARCserve Program:

- Quick Access— Introduces the system and provides quick links to ARCserve Manager functions.
- Wizards— Simplifies the most common tasks of ARCserve. You can access the Backup, Device, Restore, Disaster Recovery, Device Configuration and Central Database Configuration Wizards.
- Utilities— Offers several utilities that you can use to manage files. The utilities are Recover DB, Merge, Scan, Compare, Count, Purge, and Copy.
- News— Provides notes and the latest product information.
- Support— Links you to the tools you can use to help solve problems with your computer. The tools include Activity Log, E Support, Certified Device List and E News.

Managers

ARCserve provides quick links to ARCserve management functions. You can access these functions by selecting Quick Access on the Home Page or from the Quick Access link in the left pane. The Managers include:



Backup Manager

Backs up data to media. You can schedule and submit jobs to back up your machines. Information about each backup job (such as the path and name of each file, as well as the media used) is logged in the ARCserve database. Client agent software for other operating systems (i.e., OS/2, NetWare, MAC, UNIX, Windows 95 and 3.x) and Database agents running under the 2000 platform are also available.



Restore Manager

Restores data that has already been backed up by ARCserve. This button brings up the Restore Manager so that you can schedule and submit jobs to restore the data from your machines. It allows you to perform a complete or partial restore of data.



Database Manager

Displays information from the ARCserve database, such as the jobs processed by ARCserve, the media used by ARCserve, and the devices you are using with ARCserve.



Job Status

Monitors all pending, completed, and active jobs from the Job Status Manager window. Reschedule pending or completed jobs. Submit new jobs, delete jobs, and stop active jobs. Log information is provided for each completed job.



Devices Manager

Displays information about your storage devices and media. It also allows you to change a drive's compression mode, and perform media functions such as compression, formatting, erasing, ejecting and retensioning. ARCserve supports a wide variety of media including 4mm, 8mm, DLT, QIC, removable optical drives, Iomega's Zip or Jazz media, PDs, and MO formats.



Report Manager

Generates reports from data in the ARCserve database. Various reports include information about jobs, media errors, sessions, backup devices, and media pools.



Media Pool

Manages, creates and maintains logical groupings of media for easy identification of backups, which allow efficient scheduling of the maintenance and recycling of your media. You can design specific media rotation schemes to suit the particular archive needs. Reports detail the media information for media pools used in your rotation schemes, including media status and session information.



ARCserve
Server Admin

Allows you to manage the three ARCserve Engines or Services: Job, Tape and Database.



Alert Manager

Sends messages to people in your organization using various methods of communication, regarding events that occur during the functioning of ARCserve.

Wizards

ARCserve also provides you with wizards to simplify the most common tasks of ARCserve. You can access the Wizards by selecting Wizards on the Home Page or from the Wizards link in the left pane. The Wizards include:



Backup
Wizard

Takes you step-by-step from initial target selection to scheduling a repeating job.



Device
Wizard

Simplifies the task of formatting, erasing, re-tensioning, compressing, and ejecting the storage media.



Restore
Wizard

Restores data from media to a specified destination. You can restore by session or by querying for the specific data.



Disaster
Recovery-
Boot Disk
Wizard

Creates local or remote Disaster Recovery Boot Disk Set using this step-by-step tool.

Note: The Disaster Recovery Boot Disk Wizard is only available if the Disaster Recovery Option, which is sold separately, is installed.



Device
Configuration

Allows you to configure the backup devices and set the RAID levels on your system using this step-by-step tool.



Central
Database
Configuration

Allows you to configure the relationships between the central database and the databases of the other ARCserve machines in the domain.

Utilities

ARCserve offers several utilities that you can use to manage files. You can access the Utilities from the Utilities link in the left pane. These utilities include:



Recover DB

The Recover Database utility allows you to recover the ARCserve database by scanning media containing data that was stored by ARCserve.



Merge

ARCserve creates a catalog file of your media to improve performance for merge jobs and database backup jobs. The Merge utility allows you to merge information from ARCserve media into the ARCserve database.



Scan

The Scan utility allows you to scan ARCserve media for information about your previous backups.



Compare

The Compare utility allows you to compare the contents of ARCserve media to the files on your machine.



Count

The Count utility allows you to count the files and directories on a machine.



Purge

The Purge utility allows you to purge (or delete) files from a machine.



Copy

The Copy utility allows you to copy (or move) files from one hard location to another.

Each utility is described in more detail in the following sections.

Recover DB

Each time you run a backup job, ARCserve records information in its databases about the machines, directories, and files that have been backed up, and the media that were used. This allows ARCserve to locate files whenever you need to restore them. This database information is backed up whenever you back up your ARCserve home directory.

The Recover Database utility lets you restore this information to your ARCserve database. Unlike the Merge utility, which appends database information from media to your existing database files, the Recover Database utility overwrites existing database files. You could use the Recover Database utility any time you lose your ARCserve database files, such as after a disaster like a fire, flood, or machine crash. ARCserve provides several types of advanced database recovery options:

- Operation - Specify a date to restore the database. The database session created on or before this date will be used.
- Backup Media - Specify media options for the job, such as the media timeout period.
- Pre/post - Run commands or batch files before the job runs and/or after it finishes.
- Log - Determines the level of detail you want recorded into the Job Queue's Log.
- Alert - Send messages about events in your operation.

Merge

Appends database information to your existing database files. Use the Merge option when you want to restore data from a different server than the one where your ARCserve database resides. For example, if a backup was created using ARCserve on a different machine, you can use Merge to get the media information into the database in ARCserve's home directory. Merge allows you to take a media that contains one or more backup sessions and merge the information from that media into your ARCserve database. ARCserve provides several types of advanced merge options:

- Backup Media - Specify media options for the job such as the media timeout period
- Pre/post - Run commands or batch files before the job runs and/or after it finishes
- Log - Determines the level of detail you want recorded into the Job Queue's Log
- Alert - Send messages about events in your operation

Scan

Gives you information about your media backup sessions. Each source that you select to back up is saved on media as an individual session. You can scan a single session or an entire media. Results of the Scan job can be seen in the Job Queue. You would need to do this if you are trying to recover an ARCserve machine and you need to find the most recent backup of the ARCserve database so that you can restore it. You could also use the Scan utility if you want a list of the files that were backed up. ARCserve provides several types of advanced scan options:

- Backup Media - Specify media options for the job, such as the media timeout period
- Operation - Specify some general options for the job such as to scan files with CRC verification, or to enable database recording.
- Pre/post - Run commands or batch files before the job runs and/or after it finishes.
- Log - Determines the level of detail you want recorded into the Job Queue's Log.
- Alert - Send messages about events in your operation.

Compare

Compare the contents of a media session to files on a machine. Results of the Compare job can be seen in the Job Queue. You could use this option after a backup to verify that the backup copied all of the files to media without error. ARCserve provides several types of advanced compare options:

- Backup Media Specify media options for the job such as the media timeout period.
- Operation - Specify whether or not to enable database recording.
- Pre/post - Run commands or batch files before the job runs and/or after it finishes.
- Log - Determines the level of detail you want recorded into the Job Queue's Log.
- Alert - Send messages about events in your operation.

Count

Count the number of files and directories on a machine. Results of the Count job can be seen in the Job Queue. You could use this option after a Copy job to verify that the Copy function copied all of the files from one disk to another without error. ARCserve provides several types of advanced count options:

- Operation - Specify whether or not to enable database recording.
- Pre/post - Run commands or batch files before the job runs and/or after it finishes.
- Log - Determines the level of detail you want recorded into the Job Queue's Log.
- Virus - Scan files for viruses before they are counted.
- Alert - Send messages about events in your operation.

Purge

Allows you to delete files and directories from a machine. Results can be seen in the Job Queue. ARCserve provides several types of advanced purge options:

- Operation - Specify some general options for the job such as to remove directories or enable database recording.
- Pre/post - Run commands or batch files before or after the job runs.
- Log - Determines the level of detail to record in the Job Queue's Log.
- Alert - Send messages about events in your operation.

Copy

Allows you to copy files from one disk to another. For example, you can run a copy job on your local machine to store files and directories on another machine that is going to be backed up to media. Copy options determine related actions that occur during or after the copy operation:

- **Retry** – Specify when to retry open files and file sharing options.
- **Operation** – Specify the following options as well as whether or not to enable database recording.
 - Mirror** – Create a mirror image of the source machine on the destination. Copy files and directories from source and delete all files and directories that do not appear on source.
 - Delete Files** – Delete files from source machine after copying to destination.
 - Copy File Security Information** – You can copy the file-level security information for file access on NTFS volumes.
- **Destination** – Specify options for the directory structure and file conflict resolutions.
- **Pre/Post** – Run commands or batch files before the job runs and/or after it completes.
- **Log** – Determines the level of detail you want recorded into the Job Queue's log.
- **Virus** – Scans files for viruses before they are copied.
- **Alert** – Send messages about events in your operation.

Preferences

ARCserve allows you to set preferences. The available tabs are:

Global Settings

Set Refresh Rate

- **Set Job Queue Refresh Rate** – Specify a time, in seconds, for periodic update of the Job Status Manager.
- **Set Device Manager Refresh Rate** – Specify a time for periodic update of the Device Manager.
- **Set Animated Speed** – Specify a speed at which the tape bitmap will rotate if animation is selected for the Device or Backup Manager.

Enable/Disable Options

- Show Removable Drive in Tree – This will display removable drives within the Tree view. For example, this will display a floppy drive under My Computer.
- Show CDROM Drive in Tree – This will display CD ROM drives under My Computer within the tree view.
- Show Leaf Node in Source Tree – This will display all leaf nodes within the tree view. This means that files will be displayed under directories, and that media will be displayed beneath drives
- Auto Start All Engines – All ARCserve engines will start automatically when the Manager is started.

Prune Activity Log Settings

Check the Enable Auto-Pruning of the Activity Log box to allow ARCserve to enables the following two options:

- Prune Entire Activity Log – Running the prune job will remove the entire activity log
- Prune Activity Log that is Older than [x] [time units] – Select this to remove activity log entries older than a certain time period (default is all entries older than one month). Time units include Days, months, and years.

Library Filter

Note: This option applies to library devices. It allows the user to select the active view.

- Show Empty Slots – Empty slots in the tape library will be displayed when this box is checked.
- Show slots between X and Y – Use this to only show a selected portion of all slots within the tape library.

Backing Up Your Data

ARCserve allows you to back up most machines attached to your Windows NT/2000 network, using one of the following sources:

- the administrative shared drives
- user-shared files, directories, and drives

User-shared drives and directories are displayed in ARCserve's Browser, so that you can quickly select them for backup. Since ARCserve separates and lists Windows 2000 machines by the domain or workgroup to which they belong, you can easily back up all the machines belonging to a specific domain or workgroup, by selecting the name of the domain or workgroup.

The optional ARCserve Client Agents allow you to communicate with remote workstations in various environments. This provides complete system backups, including system information of non-Windows 2000 systems, such as NetWare, Macintosh, UNIX, or OS/2.

Similarly, the optional Backup Agents allow ARCserve to back up and restore online databases and applications such as MS Exchange, Lotus Notes, MS SQL Server, Oracle, and Informix.

For more information on ARCserve's Client Agent and Backup Agent options, refer to the *ARCserve Getting Started* Guide or contact your Computer Associates representative.

The Backup Wizard

ARCserve's Backup Wizard guides you step-by-step through the process of creating and submitting a backup job of a single machine without running the ARCserve Backup Manager. The Backup Manager is provided to allow you to perform more customized backups using filters, options and scheduling.

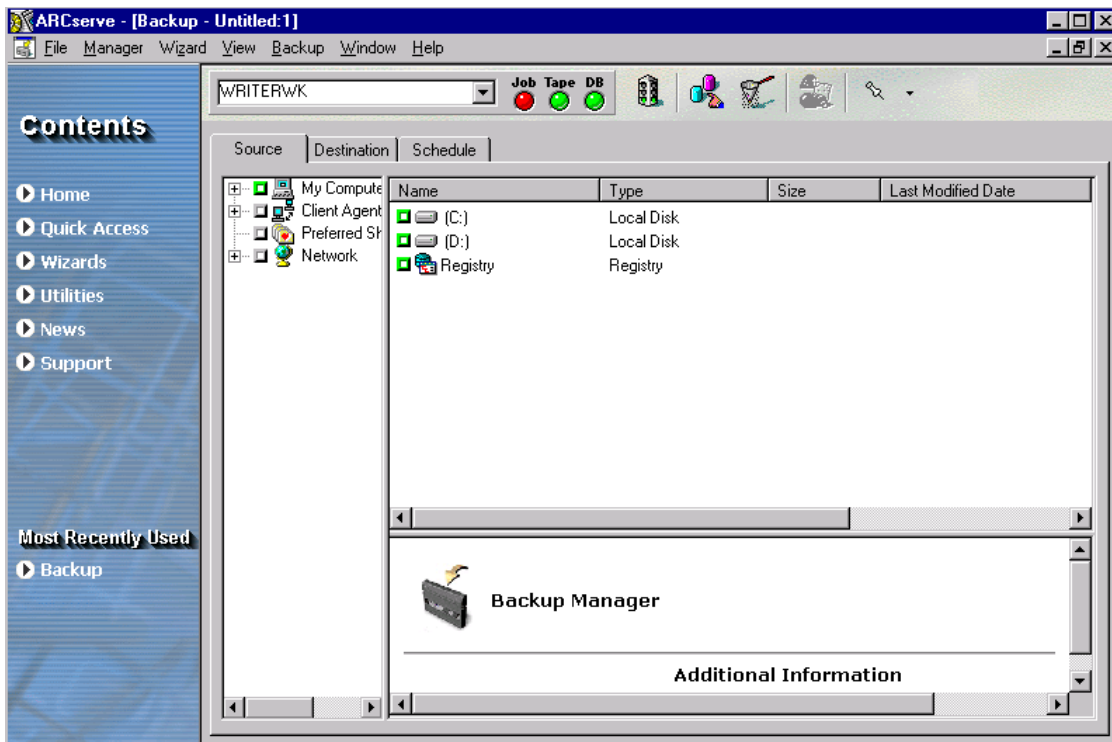
The Backup Manager

For more customized backups, you can use Backup Manager to:

- Back up to various media or create a customized backup scheme.
- Use filters to selectively exclude or include directories and files from backup jobs.
- Create an automated backup scheme using ARCserve's GFS rotation scheme.
- Apply backup options to local source objects (such as volumes and nodes) or globally to the entire backup job, or to both at the same time.

ARCserve allows you to back up (either fully or partially) the Windows 2000 registry for a Windows 2000 machine as well as the system state. Each backup job requires a source and a destination (media). The Backup Manager screen contains three tabs to customize your backup job:

- Source
- Destination
- Schedule



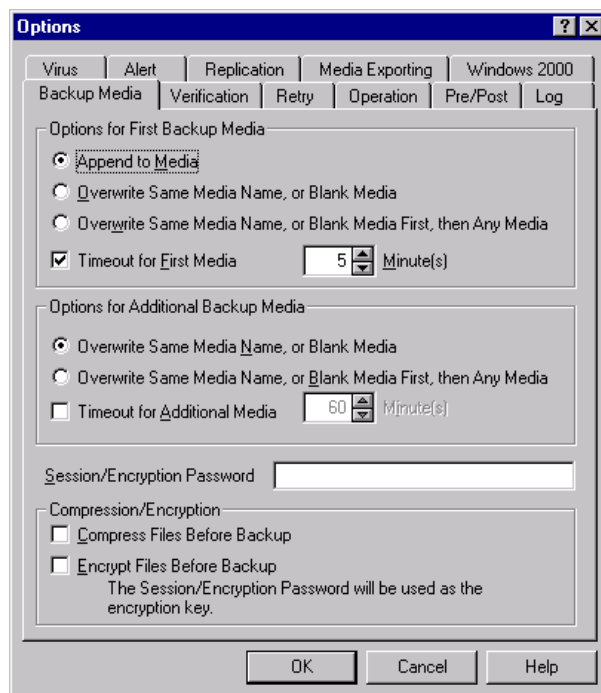
Backup Options

This section describes some of the basic backup options available to you. For a description of additional backup job options, filtering, and job scheduling features, refer to “Customizing Your Jobs” in this guide. The available option tabs are:

- Backup Media
- Verification
- Retry
- Operation
- Pre/Post
- Log
- Virus
- Alert
- Replication
- Media Exporting
- Windows 2000

Backup Media

You can specify the overwrite/append rules for the media used in your backup job while you are configuring the job. This section describes the rules so that you can determine which method is best for your purposes.



Options for First Backup Media

The First Backup Media is the media you use when the backup job begins. The options for the first media are:

- **Append to Media**— Adds job sessions to the selected media.
- **Overwrite Same Media Name, or Blank Media**— Overwrites the media in the drive only if it is the name you specified for the job or if the media is blank. If neither of these conditions is met, ARCserve prompts you to supply the specific media name.
- **Overwrite Same Media Name or Blank Media First, then Any Media**— Overwrites the media in the drive only if it is the name you specified for the job or if the media is blank. If neither of these conditions is met, ARCserve reformats whatever media it finds in the drive and starts backing up files at the beginning of the media.
- **Timeout for First Media**— Specifies the number of minutes ARCserve will wait for a media to be inserted into a drive before it cancels the job.

Options for Additional Media

The additional backup media rules apply to jobs that require more than one media. You need to specify which media ARCserve can use when the job spans media. The additional backup media rules are:

- **Overwrite Same Media Name, or Blank Media** – Writes to the media in the device only if it has the same media name (but a different media ID) or if it is blank. When the job requires an additional media, ARCserve checks if the new media has the same name (but different media ID) or if it is a blank media. As long as the ID is different, ARCserve reformats the media, giving it the same name and ID as the first media. Only the sequence number is different.
- **Overwrite Same Media Name or Blank Media First, then Any Media** – Overwrite any media found in the device (as long as it has a different ID from the first media's ID). All subsequent media are reformatted with the same name and ID as the first media. Only the sequence number will be different.
- **Timeout for Additional Media** – Specifies the number of minutes ARCserve will wait for a media to be inserted into a drive before it cancels the job.

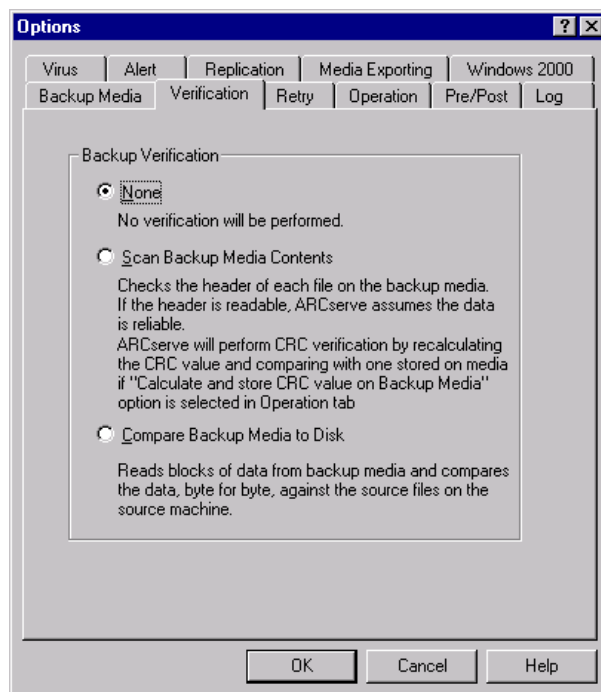
Secure Your Backup Data

To ensure security throughout your network, you can choose from the following methods:

- **Session/Encryption Password** – To enhance network security, specify a session/encryption password that will be required in order to restore this data from media.
Note: Make sure that you remember your password so that it is available when you want to restore your data. If you do not know your password, you will not be able to access the data!
- **Compress Files Before Backup** – Specify whether or not files will be compressed before backup. You can increase the effectiveness of your media and increase data speed by setting this option on. If the tape drive does not support compression, or if the hardware compression is disabled, ARCserve will compress the data. On a compressed NTFS volume, the data being archived will be stored compressed.
- **Encrypt Files Before Backup** – Specify whether or not files will be encrypted before backup. Encryption encodes your data so that it is not intelligible.

Verification

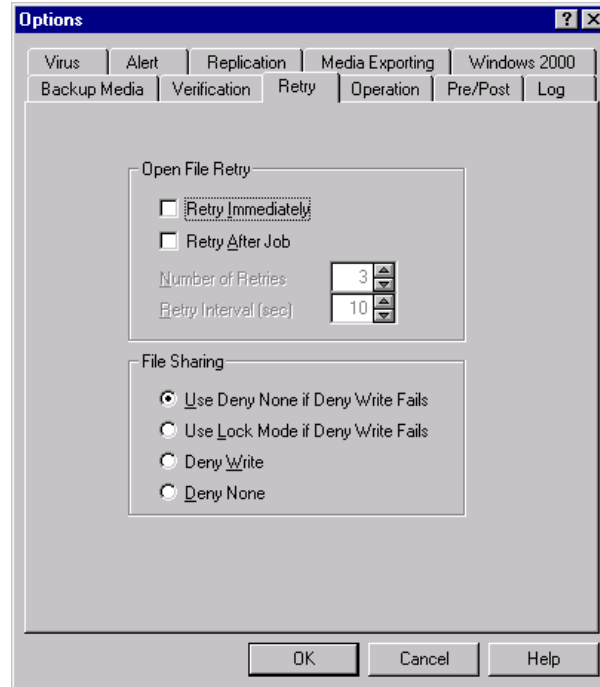
ARCserve allows you to verify that your data was correctly backed up to media. You can verify data for the entire backup job or for a selected drive in your backup job. The global verification options (applied to the entire job) will be overridden by the options selected for a drive.



- None – The back up will not be verified.
- Scan Backup Media Contents – Check the proprietary ARCserve data area (the header) of each file on the backup media. If it is readable, ARCserve assumes the data is reliable. If it is not readable, the Activity Log is updated with this information. This is the fastest verification method.
- Compare Backup Media to Disk – Data from the backup media is read and compared byte for byte against the source files. This option takes time, but ensures that all data on the backup media are exactly as on the disk. If ARCserve finds a mismatch, the errors are recorded in the Activity Log.
- CRC (Cyclic Redundancy Check) – If you selected this option on the Operation tab, ARCserve will automatically perform CRC verification. This method assigns a value to the data that you copied to media and compares it to the value assigned to the data that you backed up. This enables you to identify the individual data packets that were backed up.

Retry

When ARCserve encounters an open file, there are two types of Retry options:



- Open File Retry – determines how frequently (if at all) ARCserve will attempt to back up or copy open files.

Retry Immediately – Back up or copy the file again, immediately after the first attempt failed. If the file is still unavailable, this information is written to ARCserve’s Activity Log, and the job is labeled “Incomplete”.

Retry After Job – Back up or copy the file again after all the other source files have been backed up. If the file is still unavailable after these attempts, this information is written to ARCserve’s Activity Log, and the job is labeled “Incomplete”.

Number of Retries – Number of times you want ARCserve to try to back up or copy the file.

Retry Interval (sec) – Period of time you want ARCserve to wait between attempts.

- File Sharing – determines how ARCserve shares the files with other applications when backing up or copying a file.

Use Deny None if Deny Write Fails – [default] Attempt to place the file in “Deny Write” mode. If this is not possible (because the file is already open), then place the file into “Deny None” mode.

Use Lock Mode if Deny Write Fails— Attempt to place the file in “Deny Write” mode. If this is not possible (because the file is already open), then lock the file completely (prohibiting any user from opening or writing to the file). This option ensures the most recent version of the file is backed up or copied.

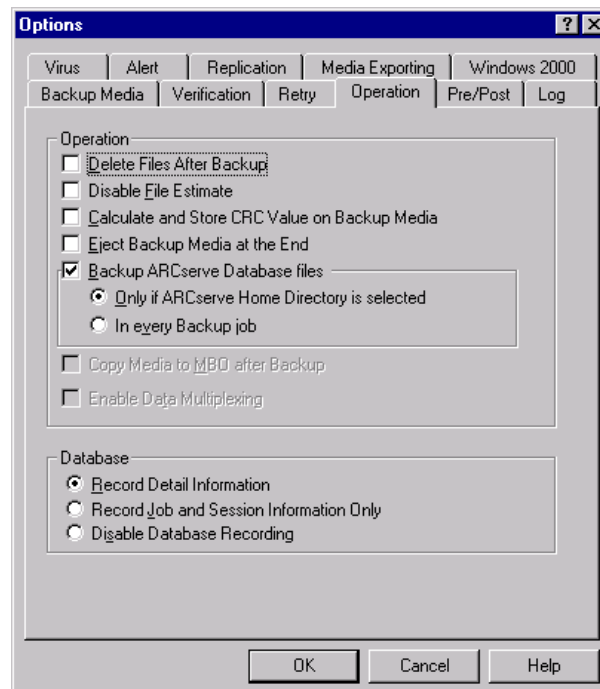
Deny Write— Prevent another process from writing to the file while ARCserve has it open. If another process has the file open before ARCserve can open it, ARCserve will not back up the file (unless you specified an Open File Retry option).

Deny None— Allow another process to read or write to the file, regardless of whether ARCserve opens it first. This option ensures that your files are up-to-date, although the file that was backed up or copied may not be the most recent version.

Note: If you use applications (such as e-mail) that are in operation 24 hours a day, you may want to choose one of the Deny None methods. As long as no other process is writing to these files during the job, the backup or copy will be consistent. If you want to ensure that only the most current version of every file is backed up or copied, you should select a Deny Write or Lock Mode option.

Operation

The Operation options for Backup determines related actions that occur during or after the backup, and the level of information that is recorded in the database.



The Operation tab has the following options:

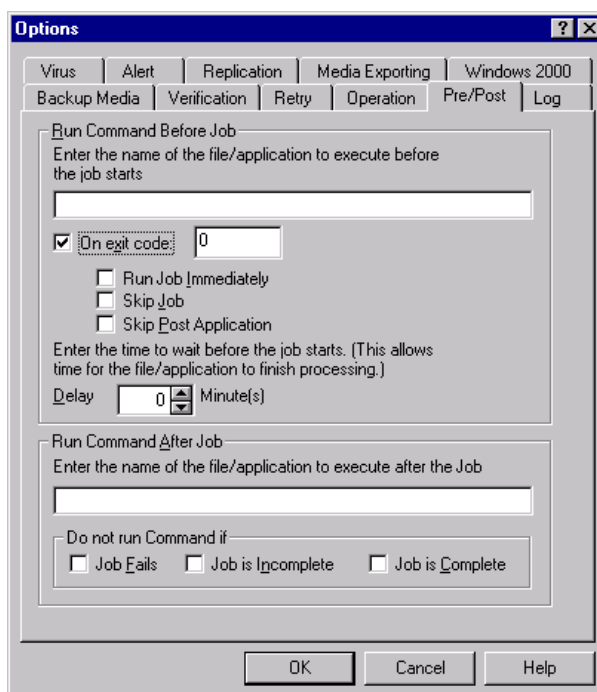
- Delete Files After Backup – Delete source files from the source machine after they have been backed up to media. Use this option to perform disk grooming. Set a filter to back up files that haven't been accessed for a certain period of time (say, three months or so), and then select this option to delete these files from the source disk.
- Disable File Estimate – By default, before any file is backed up to media, ARCserve performs an estimate of how long the job will take. Select this option if you want ARCserve to skip this function.
- Calculate and store CRC value on Backup media – Calculating and storing the CRC value on the backup media will enable ARCserve to perform CRC verification during the backup job. To instruct ARCserve to use the CRC value stored on media, see the Backup Options, Verification tab.
- Eject Backup Media at the End – Eject the media from the drive after the job finishes. This helps prevent any other job from overwriting information on this media.
- Backup ARCserve Database files – Select this option to back up the ARCserve database files with your backup job. You must also select whether to backup the database on every backup job or only when the ARCserve home directory is backed up, by clicking on the appropriate option button.
- Database Options are:
 - Record Detail Information – Log all the information pertaining to the source and destination in the database.
 - Record Job and Session Information Only – Log only job and session information in the ARCserve database.
 - Disable Database Recording – Don't log any information about this job in the database. Select this option if you are backing up the ARCserve database or if you have limited disk space.

Pre/Post

ARCserve allows you to specify pre/post commands that perform various actions immediately before and/or after your backup job:

- For the entire backup job: These commands are executed at the beginning or end of the job.
- For a selected machine: These commands are executed before and/or after the data from the selected machine is backed up.

Note: If you specify pre/post commands for the entire job, as well as for machines in the job, the global commands are executed before/after the job and local commands are executed when the selected machine is backed up.



Run Command Before Job

Enter the path to and name of the file to be executed on the machine before the job takes off.

- On Exit Code – ARCserve detects exit codes of other programs. You can specify the following options for a particular exit code:
 - Run Job Immediately – The job runs immediately if the appropriate code is selected.
 - Skip Job – The job does not run if the appropriate exit code is detected.
 - Skip Post Application – Skip any commands specified to run after the job if the appropriate code is detected.
- Delay in Minutes – Specify the delay in which ARCserve waits before running a job when the appropriate exit code is detected.

Run Command After Job

Enter the path to and name of the file to be executed on the machine after the job is completed.

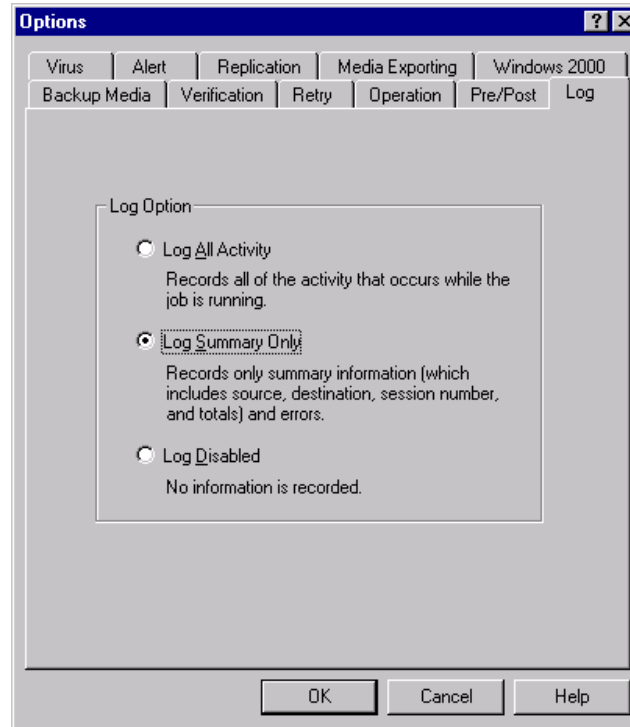
Do not run Command if

Specify for a command not to run if ARCserve detects that the following events occur:

- Job Fails – If ARCserve detects that a job fails, then a command does not run.
- Job is Incomplete – If ARCserve detects that a job is not completed, a command does not run.
- Job is Complete – If ARCserve detects that a job is completed, a command does not run.

Log

The Log options determine the level of detail that is included in the log report for the operation. The log options can be set in the following windows: Backup, Restore, Compare, Scan, Copy, Count, and Purge.

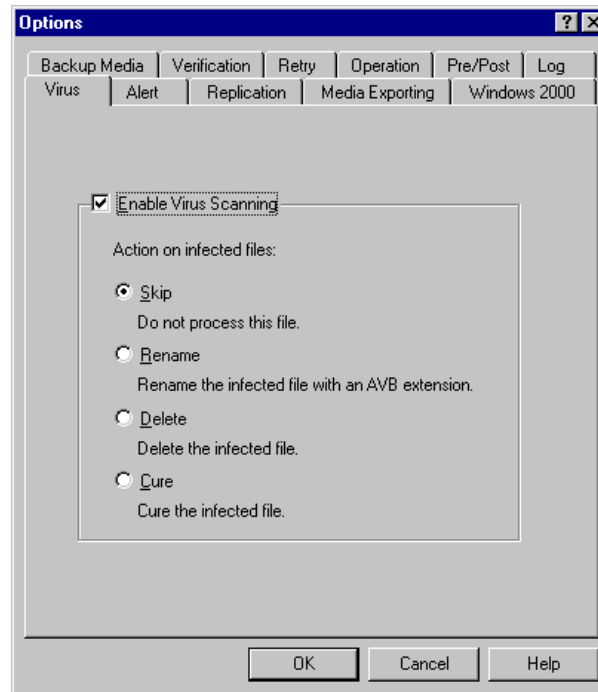


The log options are:

- Log All Activity – Record all of the activity that occurs while the job is running in the Job Log.
- Log Summary Only – Record summary information on the job (including source, destination, session number, and totals) and errors.
- Log Disabled – Do not record any information about this job in the Job Log.

Virus

You can automatically scan for viruses during a backup, copy, count, or restore operation using the virus scanning options.

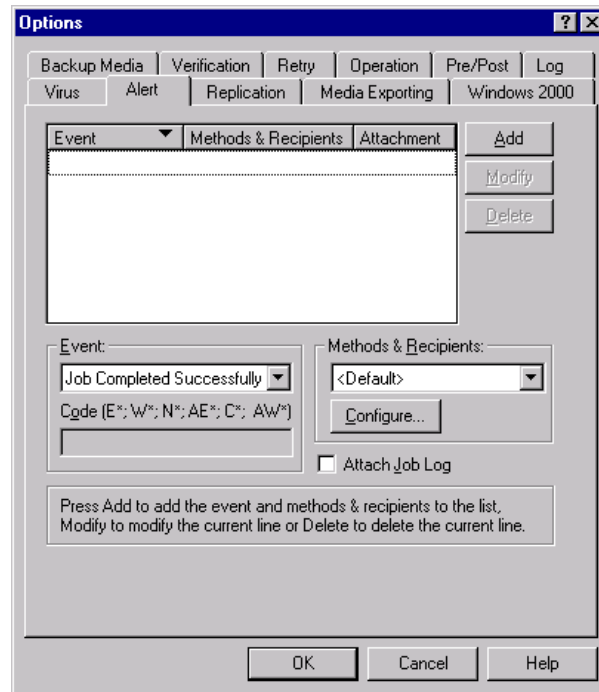


When you check Enable Virus Scanning, the available options are:

- Skip— Do not back up/restore the infected file.
- Rename— Rename the infected files with the extension AVB. (If a file with the same name and the extension AVB already exists, then extension AV1 is used, and so on.)
- Delete— Delete the infected file.
- Cure— Attempts to cure the infected file.

Alert

You can use the Alert notification system to send messages about events in your backup operation.



Choose one or more of the following events for which you want to be notified:

- Job completed successfully – All of the nodes and drives/shares were processed.
- Job incomplete – Some nodes and/or drives/shares were missed.
- Job canceled by user – The user canceled the job.
- Job failed – The job was started but could not be completed.
- Virus detected – A virus was detected in one of the files to be backed up. See Virus options (Backup, Copy, Count)
- Customized Event – A customized event occurred. To specify this type of event, enter an error, warning, or notification code in the space below the Event drop-box.

Choose one or more of the defined Alert configurations. The <default> configuration means that you will use whatever is configured in Alert Manager. Before you choose, click the Configure button to define further configurations.

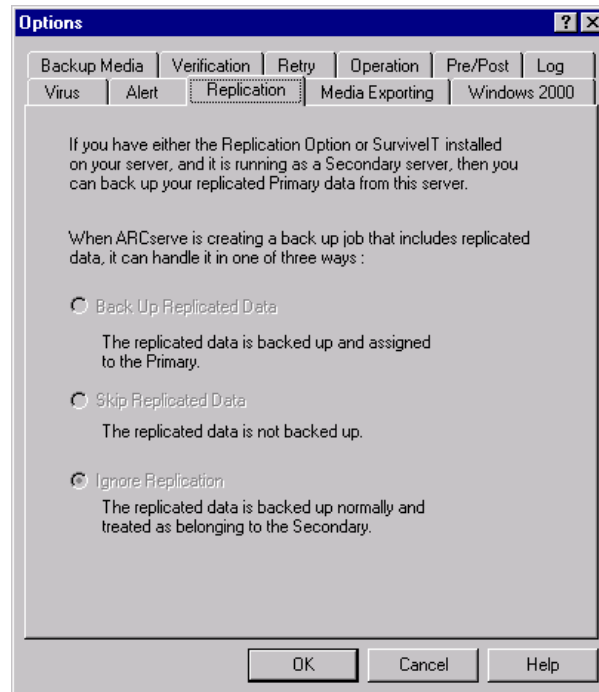
- Broadcast
- Pager
- SNMP
- Event
- Printer
- E-mail
- Lotus Notes
- Unicenter TNG

Select the *Attach Job Log* option to include the job log information in the Alert message. (This option applies for Trouble Tickets and Email only.)

Note: The list you create using Alert Options is saved with the Job Script and the configuration defined using the Configuration button.

Replication

If you have either the Replication Option or SurviveIT installed on your server, ARCserve can be set to back up replicated data on your server.

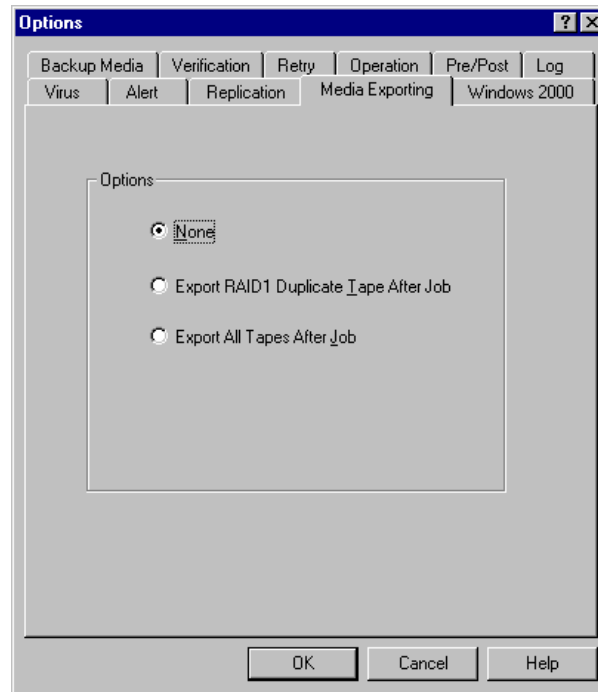


The available options are as follows:

Backup Replicated Data	ARCserve will back up the data on your server including all replicated data, which will appear to have been copied from the primary servers. Each instance of replicated data will be a separate session on the backup media. If you restore your data to the original location, the data will go to that primary machine, not to the server.
Skip Replicated Data	ARCserve will back up only the data on your server. ARCserve will skip any replicated data.
Ignore Replication	ARCserve will back up all data normally. It will not check to see if the data is replicated.

Media Exporting

You can export media at the end of a backup job. It allows you to move it out of the library or to an off site location for safe storage. If the job does verification, the export is done at the end of the verification. This feature is available for regular and rotation jobs and supported on tape and optical libraries and RAID.



The options available are:

None

No media exporting will take place at the end of a backup job.

Export Duplicate
Tape After Job

ARCserve will export all the duplicate media for the related backup. If the job spanned to multiple media, all the duplicate media used in this job will be exported.

Note: This option is for RAID 1 support for users with libraries and mail slots only.

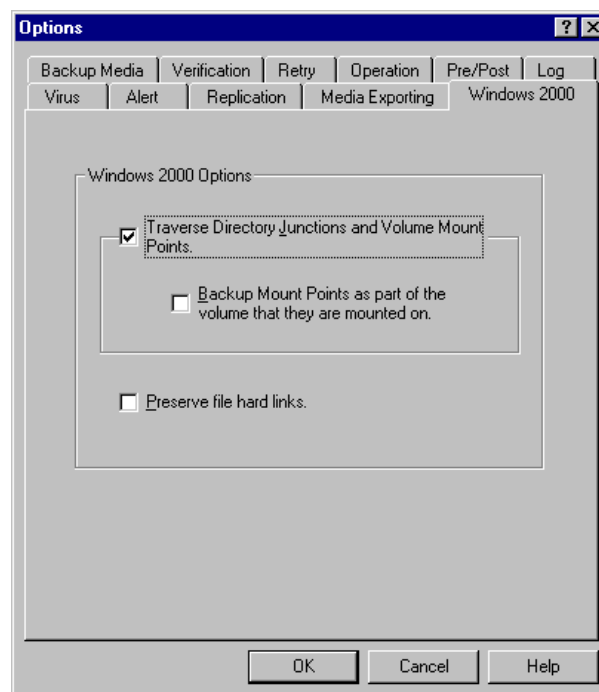
Export All Tapes After Job

ARCserve will export all the media for the related backup. If the job spanned to multiple media, all the media used in this job will be exported. If there are not enough mail slots to export all the media, the media that could not be exported will be moved back to the original home slot. In single mail slot libraries, ARCserve retries a few times to check if the mail slot is empty to move the next media to the mail slot. If the operator does not move the media, ARCserve will write this information in the activity log.

Note: This option is for RAID 1 support for users with libraries and mail slots only.

Windows 2000

You can set Windows 2000 options. This is only available for systems running Windows 2000.



- Traverse Directory Junctions and Volume Mount Points – If ARCserve encounters a volume mount point or a directory junction, it will not expand it during backup.
- Backup Mount Points as Part of the volume that they are mounted on – If this option is selected, ARCserve will backup mounted volumes and directory junctions as separate sessions.
- Preserve File Hard Links – Checking this box, forces ARCserve to preserve hard links during a restore.

Backup Agent for Open Files

If the Backup Agent for Open Files (BAOF) is installed on any of your computers, you can manage the BAOF Engine directly from the Backup Manager.

1. In the ARCserve Backup Manager, go to the Source tab and select a server from the browser in the left pane.
2. If this server has the BAOF Engine installed, right-click the computer and select the following menu items or simply click these items in the Additional Information pane, on the bottom-right of the screen:

Configure Open File Agent	This displays the BAOF Configuration screen. From here, you can set various global settings for BAOF on the selected computer.
View Open File Agent Status	This displays the BAOF Status screen. This shows which files and groups BAOF is currently processing on the selected computer.
View Open File Agent Log File	This displays the Log File Viewer screen. This shows the log file for the selected computer.

See the *Backup Agent for Open Files Administrator Guide* for further details.

Disaster Recovery

The most important thing you can do to ensure against data loss is to maintain current backups of all your servers and workstations. If you do not have these backups, then ARCserve is limited in its ability to recover against data loss or disaster. Be sure to create a media rotation scheme and scheduling to maintain current backups.

Refer to the *Getting Started* Guide for information on the Disaster Recovery Option.

The Restore Wizard

ARCserve's Restore Wizard guides you step-by-step through the process of restoring your data. With the Wizard you are able to submit a restore job to ARCserve's job queue without running the ARCserve Restore Manager. The Restore Manager is provided to allow you to perform more customized restores using filters, options and scheduling. The Wizard offers two restore methods:

- Restore By Session (the default method)
- Restore By Query

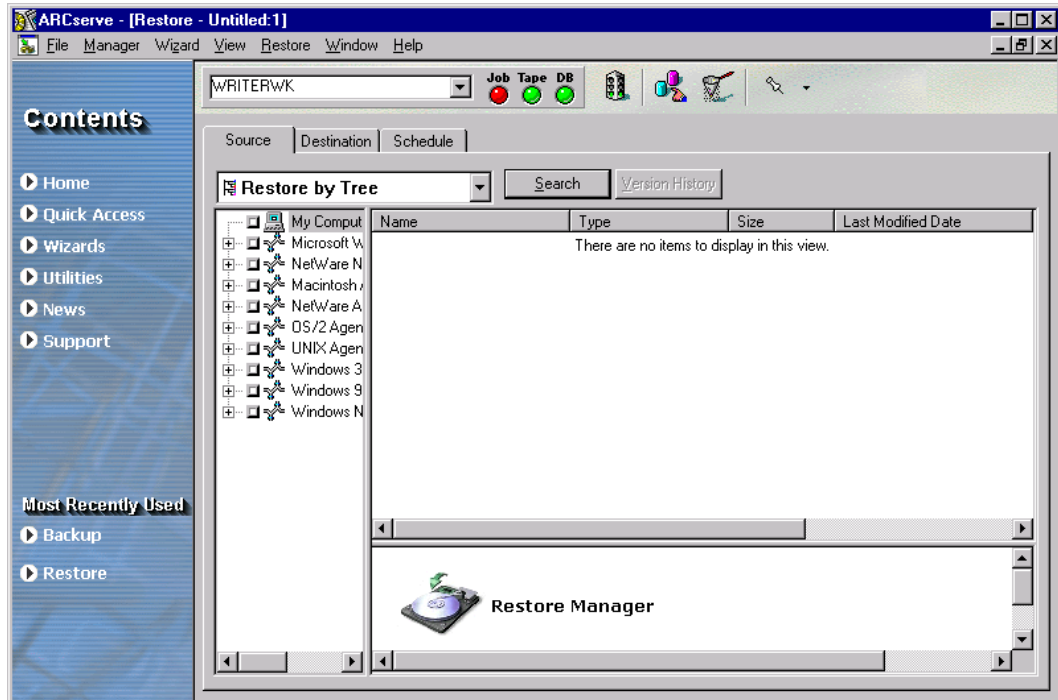
For further information on how to use the Restore Wizard, refer to ARCserve Procedures.

The Restore Manager

The aim of running a successful restore job is to quickly identify the data you need, and to retrieve it from the appropriate backup media.

ARCserve allows you to restore data to most machines attached to your Windows 2000 network. Each restore job requires a source and destination. The files selected as your source must originate from an ARCserve backup media, and the destination must be a hard drive. The Restore Manager screen contains three tabs to customize your restore job:

- Source
- Destination
- Schedule



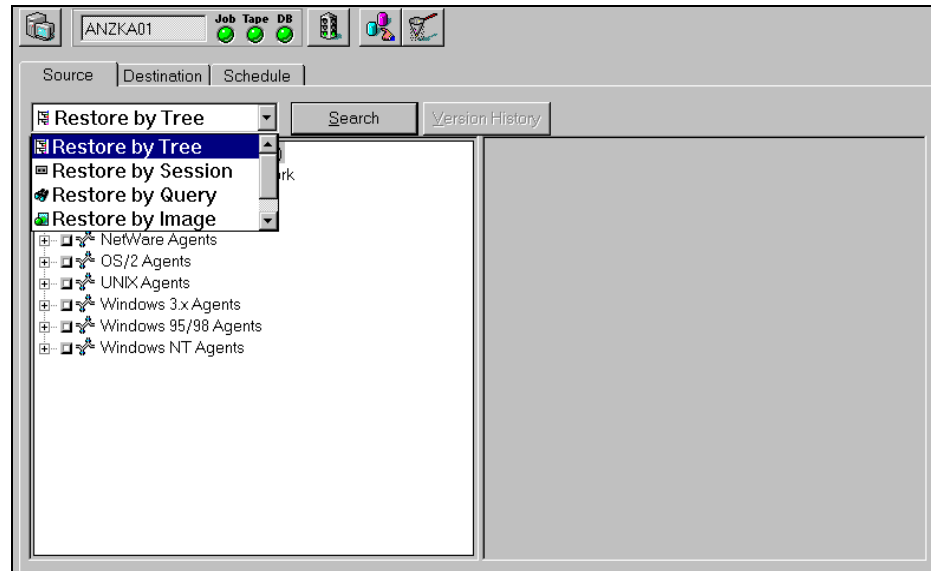
The optional ARCserve Client Agents allow you to communicate with remote workstations in various environments to restore data to non-Windows 2000 systems, such as NetWare, Macintosh, UNIX or OS/2.

Similarly, the optional Backup Agents allow ARCserve to restore online databases and applications such as MS Exchange, Lotus Notes, MS SQL Server, Oracle, and Informix.

For more information on ARCserve's Client Agent and Backup Agent options, refer to the *ARCserve Getting Started* guide or contact your Computer Associates representative. For information on how to perform a basic restore job, refer to *ARCserve Procedures*.

Source

ARCserve makes it easy to find the files you want to restore. Because your requirements and circumstances can vary, ARCserve provides you with four methods for selecting the data (the source) you want to restore:



- Restore by Tree – Restores a specific directory or drive from a display of files and directories that were backed up with ARCserve. Use this method when you don't know which tape contains the data you need, but you know which machine it came from.
- Restore by Session – Allows you to select the session, and the files and directories you want to restore. Use this method when you know the tape name, but are not certain about the session you want to restore.
- Restore by Query – Restores files based on the search pattern given 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(s) it was backed up to. This view uses ARCserve's database.
- Restore by Backup Media – Restores a complete backup session from a specified media in a storage device. All files in the session are restored to the destination, unless filters are added to the restore job. Use this method when a tape was created by a different version of ARCserve or if the database does not recognize it.

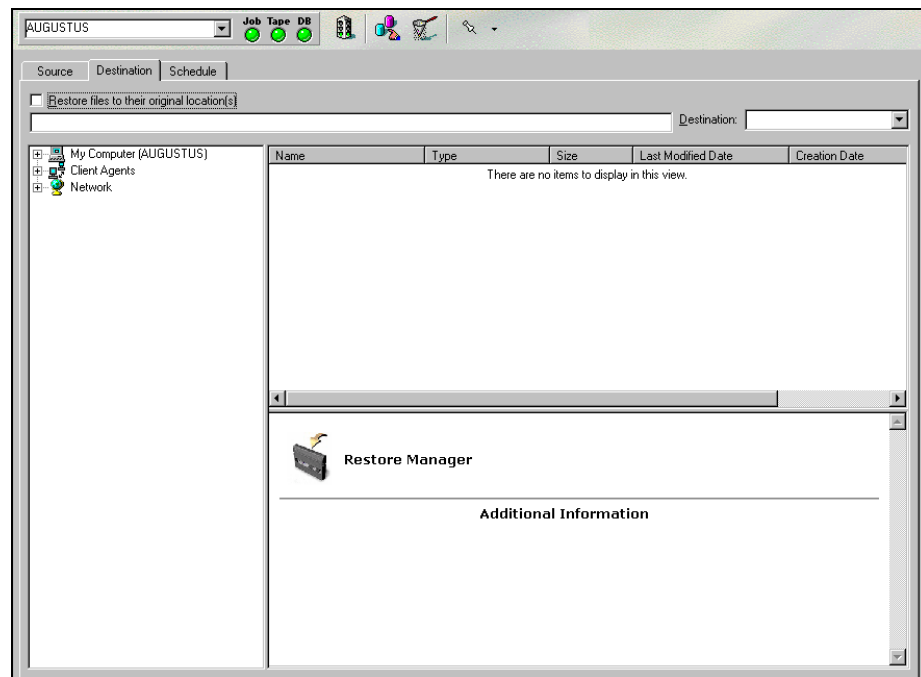
Note: If you have installed the optional Image Backup Option, you will also have Restore by Image available. Use this method when you need to read and restore blocks of data quickly by bypassing the file system.

Destination

ARCserve provides you with two methods for selecting the destination you want to restore the data to:

- Restore files to their original location
- Restore to user-shared directories and drives

Note: Only the *Restore files to their original location* method is presented by default. However, if you de-select this check box, you will be presented with a list of machines, directories and files, and you can select your specific destination from this list.



For information on how to select a destination, refer to the online *ARCserve Procedures*.


Schedule

Restore jobs can be run immediately or scheduled for a later time. Run your job immediately when:

- The job you are submitting is a “one-shot” job that you want executed immediately
- You want to monitor the job as it runs

Schedule your job when:

- The storage device is currently busy but you want to run your restore job as soon as the drive is free
- The job you are submitting will be repeated on a regular basis, whether it be daily, weekly or monthly



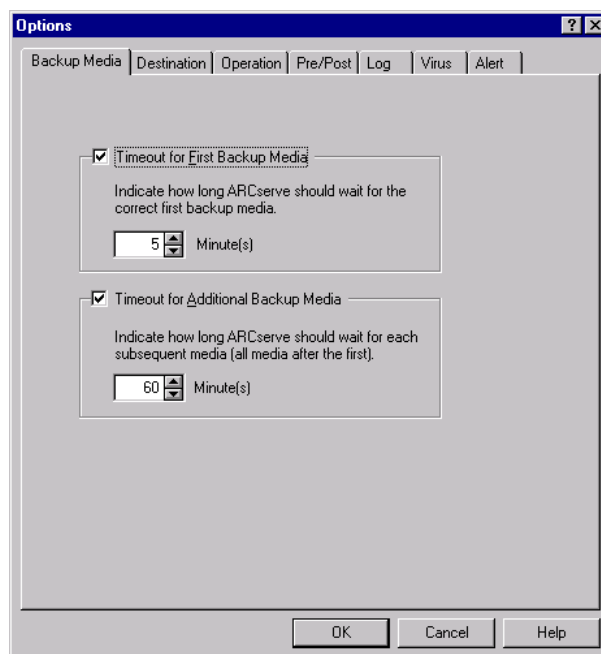
Restore Options

When you are creating a restore job, there are options to choose from to customize your job. The option tabs are:

- Backup Media
- Destination
- Operation
- Pre/Post
- Log
- Virus
- Alert

Backup Media

You can specify a timeout period that ARCserve should wait to provide the media you need to restore your data.

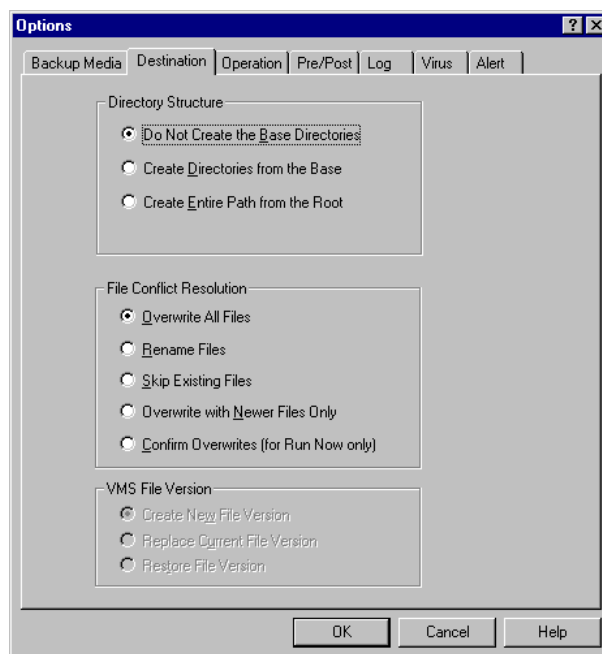


Available Media options are:

- Timeout for First Backup Media – Period of time that ARCserve waits for the first media required for your restore job. If the time expires, the job will fail.
- Timeout for Additional Backup Media – Period of time that ARCserve waits for any additional media to become available.

Destination

Determines how the directory structure is created on the destination when files are restored. They also determine which files (if any) can be overwritten.



Directory Structure Options

Select the method ARCserve should use to create directories on your destination. The Directory Structure options are:

- Do Not Create the Base Directories – (default) Do not create the base directory on the destination path, but create all subdirectories below the source base directory.
- Create Directories from the Base – Create the destination path beginning from the base directory.
- Create Entire Path from the Root – Create the entire source path on the destination. No files from any parent directories are restored. Only the directory path to the base directory is created on the destination.

File Conflict Resolution Options

Select the method ARCserve should use when there are files on the destination disk that have the same name as files being copied from the source:

- Overwrite All Files – Restore all source files to the destination regardless of conflicting file names. The files from the source will overwrite existing files on the destination.
- Rename Files – Copy the source file to the destination with the same filename but a different extension. The extension will maintain the first two characters of the original, but the last character will be 1, 2, 3..., depending on how many files ARCserve has encountered with the same name.

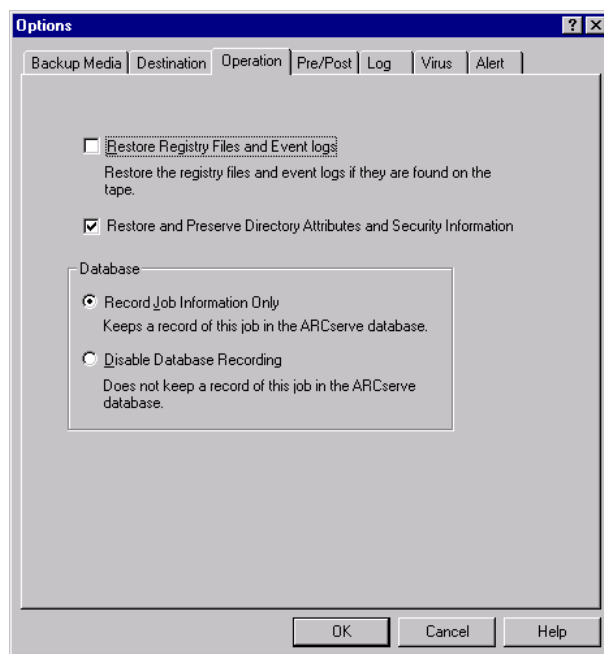
- Skip Existing Files – Do not restore a source file if a file with the same name already exists on the destination.
- Overwrite with Newer Files Only – Only restore source files whose modification date is later than the modification date of the file with the same name on the destination.
- Confirm Overwrites (for Run Now only) – Before ARCserve tries to restore the source file, you will be prompted to confirm that you want the file on the destination to be overwritten when submitting a Run Now job.

VMS File Version

Note: This option is only available with the Open VMS Agent for the Enterprise Edition.

Operation

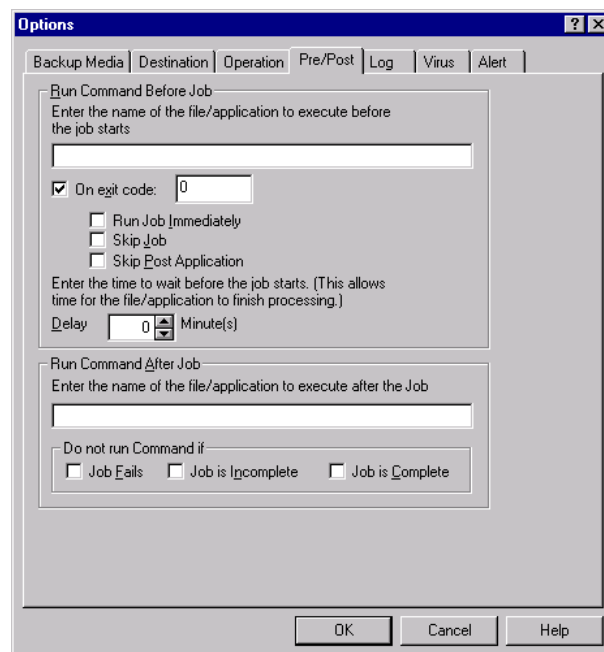
Determines related actions that occur during or after the job, and the level of information that is recorded in the ARCserve database.



- Restore Registry Files
- Restore and Preserve Directory Attributes and Security Information
- Database Recording Options are Record Job Information Only and Disable Database Recording

Pre/Post

Allows you to run a command on your machine before and/or after the job is executed. For example, you may want to use this option to load a virus scanning application before files are restored and then run a batch file you created that sends a detailed report to the printer after the job completes.



The available pre/post options are:

Run Command Before Job

Enter the path to and name of the file to be executed on the machine before the job takes off.

- **On Exit Code**— ARCserve detects exit codes of other programs. You can select to Run Job Immediately, Skip Job, or Skip Post Application for a specified exit code.
- **Delay in Minutes**— Specify the delay that ARCserve waits before running a job when the specified exit code is detected.

Run Command After Job

Enter the path to and name of the file to be executed on the machine after the job is completed.

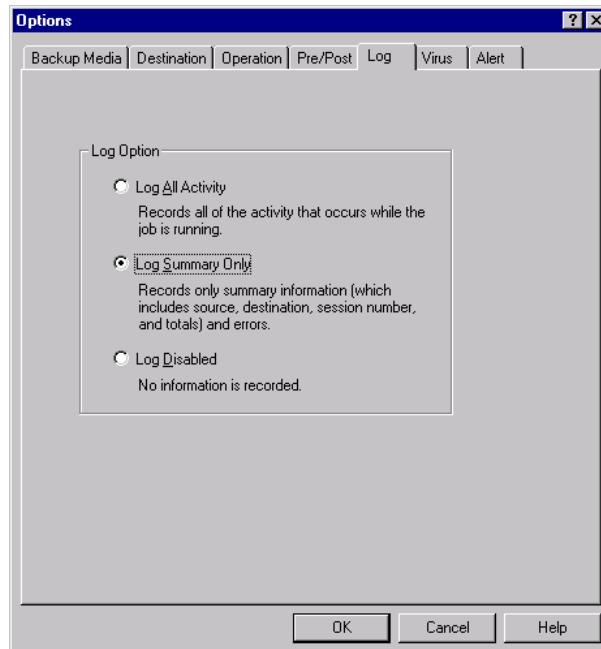
Do not run Command if

Specify for a command not to run if ARCserve detects that the following events occur:

- **Job Fails**— If a job fails, then the command will not run.
- **Job is Incomplete**— If a job is not completed, then the command will not run.
- **Job is Complete**— If a job is completed, then the command will run.

Log

Determines the level of detail that is included in the log report for the restore job. You can view the log report in the Job Queue or Database Manager window (Job View).

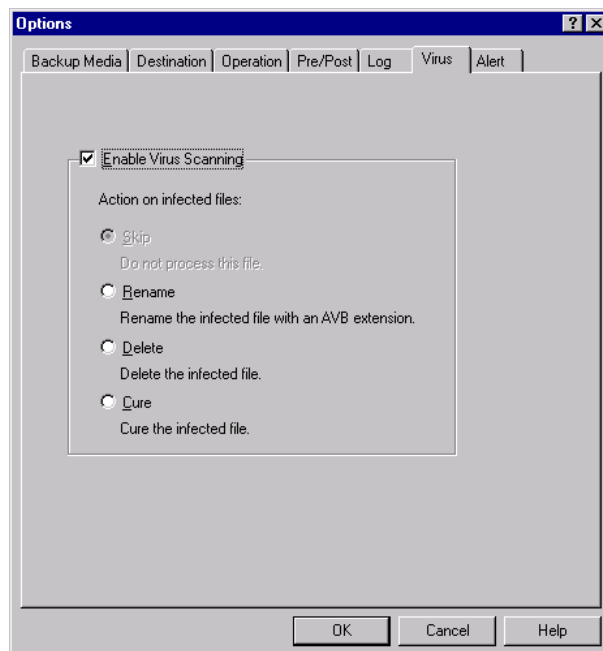


The log options are:

- Log All Activity – Record all of the activity that occurs while the job is running.
- Log Summary Only – Record summary information of the job (including source, destination, session number, and totals) and errors.
- Log Disabled – Do not record any information about this job in the Job Log.

Virus

Allows you to automatically scan for viruses during the restore operation.

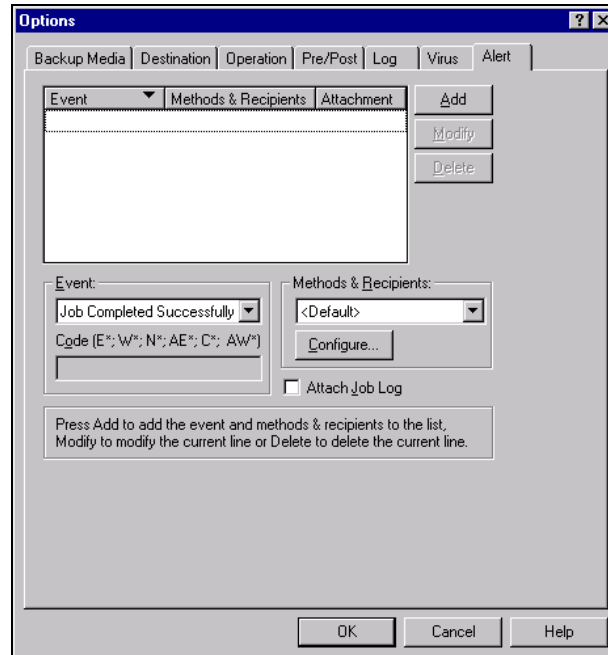


When you Enable Virus Scanning, the available options are:

- Skip – Do not back up/restore the infected file.
- Rename – Rename the infected files with the extension AVB. (If a file with the same name and the extension AVB exists, then extension AV1 is used, and so on.)
- Delete – Delete the infected file.
- Cure – Attempts to cure the infected file.

Alert

Sends messages about events in your restore operation.



Choose one or more of the following events for which you want to be notified:

- Job completed successfully – All of the nodes and drives/shares were processed.
- Job incomplete – Some nodes and/or drives/shares were missed.
- Job canceled by user – The user canceled the job.
- Job failed – The job was started but could not be completed.
- Virus detected – A virus was detected in one of the files to be backed up. See Virus options (Backup, Copy, Count)
- Customized Event – A customized event occurred. To specify this type of event, enter an error, warning, or notification code in the space below the Event drop-box.

Choose one or more of the defined Alert configurations. The <default> configuration means that you will use whatever is configured in Alert Manager. Before you choose, click the Configure button to define further configurations.

- Broadcast
- Pager
- SNMP
- Event
- Printer
- E-mail
- Lotus Notes
- Unicenter TNG

Select the *Attach Job Log* option to include the job log information in the Alert message.

Note: The list you create using Alert Options is saved with the Job Script and the configuration defined using the Configuration button.

System State Restore Options

When restoring a system state session on the Windows 2000 platform, right mouse-click in order to obtain this context menu option. The following options are available in this dialog:

- Make the Restored Copy of the Active Directory Authoritative – This forces the restored copy to become the “authoritative” version of Active Directory on the system.
- Stop the Cluster if necessary to Restore the Cluster Database – This gives permission to stop a cluster service in order to restore the cluster database. Note that this will be required only for advanced servers.
- Enable Quorum Drive Selection when Quorum Location Change (with drive letter selection) – This allows the user to set the drive of the quorum resource that a cluster currently uses.
- Do not Stop World Wide Web Service – This specifies not stop the www service while the certificate server is being restored.

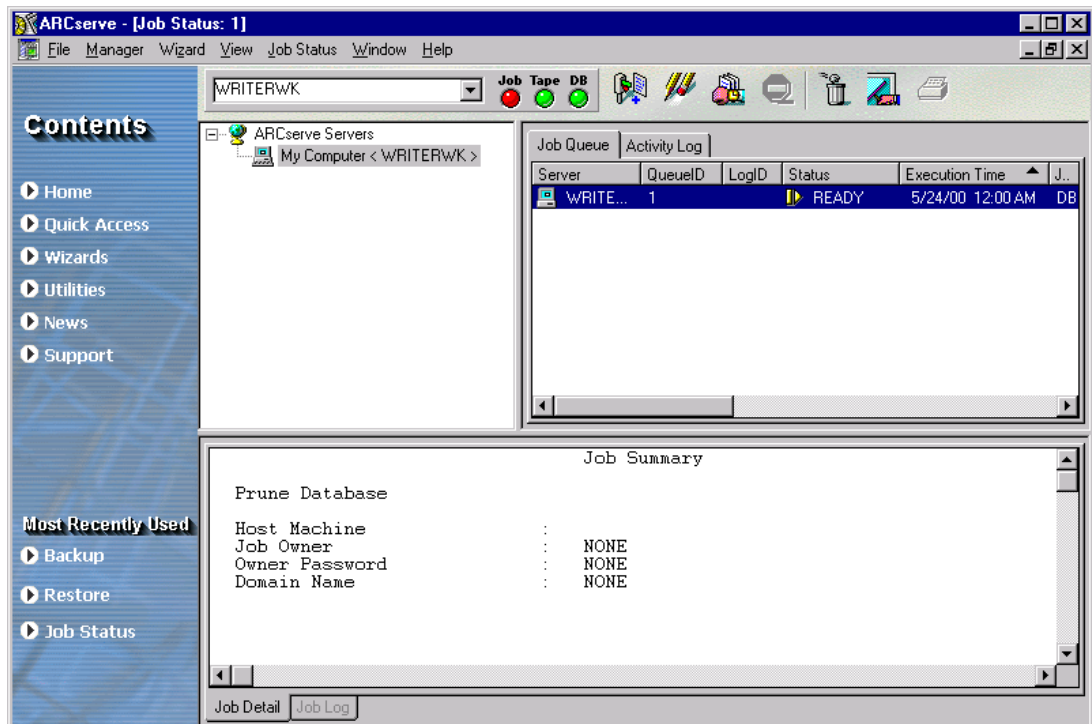
The Job Status Manager

The Job Status Manager is a graphical tool that helps you centrally manage ARCserve Servers enterprise-wide. You can use the Job Status Manager to:

- View all available ARCserve servers, job queues and activity logs.
- Manage jobs – add, delete, modify and reschedule jobs.
- Monitor progress of active jobs – View real-time status of active jobs in the queue.
- View job detail and job log information about all the jobs that have been executed.
- View activity and tape logs.
- Stop an active job.
- Modify a User Name.

Pop-up menus are available to perform various operations with the Job Status Manager. These menus appear in both the left (Server Browser) and right (Job Queue) panels. You can access a menu by clicking the right mouse button on a selection.

Tip: When accessing a pop-up menu in the Browser, the pointer must be on a selected group, server, or object. When accessing a pop-up menu in the Job Queue, a job must be selected.



Manage Servers and Groups

The left-panel browser makes it easy to navigate through ARCserve servers and their objects. You can manage your servers in groups. Select the server group and right-click on it. A menu with the following options appears:

- **New Server Group** – Appends a new server group at the end of the list and then allows you to edit the group name.
 - Rename – Allows you to edit the group name.
 - Add Member – Adds a server to the selected group.
 - Delete – Deletes the currently selected group.
- **Central Database Servers** – Contains tape session information for all of the ARCserve machines in your enterprise. You can restore the data and use the tape and session detail for generating ad hoc reports for the central database members. Administrators can use the central database to perform the following functions:
 - Query for historical information from multiple backup machines from a central location.
 - Generate ad hoc reports for all of the individual ARCserve databases in your enterprise since all of your data resides in a central location.
 - Restore machines from a central location in your enterprise.

When you select a server, a menu with the following options appears:

- Start all engines – Starts the ARCserve engines
- Stop all engines – Stops the ARCserve engines
- Delete – Deletes the server from its group
- Properties – Lists the status of the ARCserve engines for this server, build number, version, options installed.
- Security – Displays the security information for each machine and/or share that is selected as a source.

For details on using the above menu options, refer to online *ARCserve Procedures*.

Job Queue

The right panel displays information on the job queue and the activity log. Every time you run or schedule a job with the ARCserve Manager, you submit it to the ARCserve job queue. Information about all jobs (such as execution time, status, last result, and owner) is stored here. ARCserve continuously scans the job queue for jobs that are ready to be executed.

In the Job Queue panel, highlight a job and right-click the mouse button. A menu appears with the following options:

- Hold— Changes the job's status to HOLD or to READY (if it is currently on hold). (HOLD signifies that the job is not scheduled to be executed. READY means that the job can be executed.)
- Add Job— You can quickly submit a job to the queue by using a previously saved script. (A script is a job that you saved to a file. It contains the original source, destination, option, and schedule information for the job.)
- Modify Job— Modifies a job. Allows you to add options or additional sources to an existing job, without having to create a new job.
- Reschedule Job— Quickly change a job's execution date, time, or status. It also allows you to resubmit a Done job that is still in the Job Queue. You may want to do this if your job was not successful when it first ran.
- Stop Job— Cancels an active job from the ARCserve queue and reschedules it for its next regular interval.

Note: If you stop a job, the Last Result field will display "Canceled".

- Delete Job – Cancels the active job and deletes it from the ARCserve queue completely.

Important: We recommend using the Stop Job option if you want to delete an active job that repeats at intervals (determined when you create the job). Selecting the Delete Job button will interrupt and remove the job completely from the queue, and it will not be rescheduled. You will have to recreate the job if you did not save it to a script file.

- **Modify User Name**— Modify user name and password for server and source nodes.
- **Sort By**— Jobs in the queue are listed in order of execution time. This option changes the way jobs are listed in the queue. Sorting the job queue is for informational purposes only. It does not affect the order in which jobs are processed. Jobs can be sorted by clicking any of the following fields: Status, Execution Time, Job Type, Server, Last Result, Owner, Total Files, and Description.

Tip: You can resize these columns by using the “drag and drop” method with the mouse. Place the cursor on the divider between columns, click and hold down the left mouse button, and then move the divider in either direction until the column is the size you want.

- **Properties**— Double-click to call up the Job Properties dialog box when the job is pecessing.

For details on using the above menu options, refer to the online *ARCserve Procedures*.

Activity Log

The Activity Log tab contains comprehensive information about all the operations performed by ARCserve. It provides an audit trail of all ARCserve activity, including every job that is run. For each job, the log includes the following:

- 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

You can scan this log every day to see if any errors have occurred. You can also use it to find a session number in case you need to restore a specific session.

If you click your right mouse button, you have the following options:

- Fonts— Use to change the font that the log displays
- Print— Use to print the log
- Find— Use to search for specific text within the log
- Find Next F3— Use to repeat the text search
- Refresh F5— Use to update information in the log

For details on how to display the Activity Log, refer to the online *ARCserve Procedures*.

Job Detail

The bottom panel displays details about a job, including the source and destination targets and the job's schedule. If you have selected customization options such as Pre/Post backup requirements, these will be displayed as well. After a job has started, you can view its sequence and session number. If the job has completed, you can also use this panel to view the Job Log.

The Job Detail tab displays detail information on any job in the job queue. If you click your right mouse button, you have the following options:

- Fonts— Use to change the font that the log displays
- Print— Use to print the log
- Find— Use to search for specific text within the log
- Find Next F3— Use to repeat the text search
- Refresh F5— Use to update information in the log

For details on how to display the Job Detail, refer to the online *ARCserve Procedures*.

The Job Log

The Job Log displays information about specific jobs that have been executed. The level of detail in this log is determined when you set up a job. You can ask to see all of the activity that occurred while the job was running or just the summary information (which includes source, destination, session number, and totals along with any errors that may have occurred).

If you click your right mouse button, you have the following options:

- Fonts— Use to change the font that the log displays
- Print— Use to print the log
- Find— Use to search for specific text within the log
- Find Next F3— Use to repeat the text search
- Refresh F5— Use to update information in the log

For details on how to display the Job Log, refer to the online *ARCserve Procedures*.

Job Status

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

- Done— A job that has already been executed with no repeat interval.
- Ready— A new one-time or repeating job (a backup job that runs every Friday, for example) waiting to be executed.
- Active— A job that is currently being executed.
- Hold— A job that is in the queue waiting to be executed later.

Completed jobs remain listed in the Job Queue for a specified number of hours. This is set up through ARCserve's Server Admin feature. Refer to the chapter "Administering the ARCserve Server" for more information.

The Last Result field tells you whether or not your executed job was successful. If it wasn't successful, Last Result helps you determine why the job may have failed. The Last Result field may contain the following:

- Finished— All of the nodes and drives/shares were processed.
- Incomplete— The job was partially successful. Review the Activity Log information to check the exact nature of what occurred to prevent job completion.
- Canceled— The job was intentionally canceled. The following actions may have occurred:
 - A user canceled the job from the Job Queue.
 - Someone answered NO or CANCEL to a prompt.
 - The job required a confirmation of OK, or a tape needed to be inserted before the timeout was reached. (Timeout is set in the Tape options in the Backup Manager window.)

- Failed – The job failed to perform its designated task. If the job was started, but the Manager could not complete the job, you will receive “Run Failed” status (below). Review the Activity Log information to check the exact nature of what occurred to prevent job completion.
- Run Failed – The job was started, but the program that runs the job failed, because either there was not enough memory to run the job or a DLL file was not found.
- Crashed – The job was started and a system error occurred which prevented ARCserve from completing its task, such as a memory violation that caused ARCserve or the operating system to be shut down.

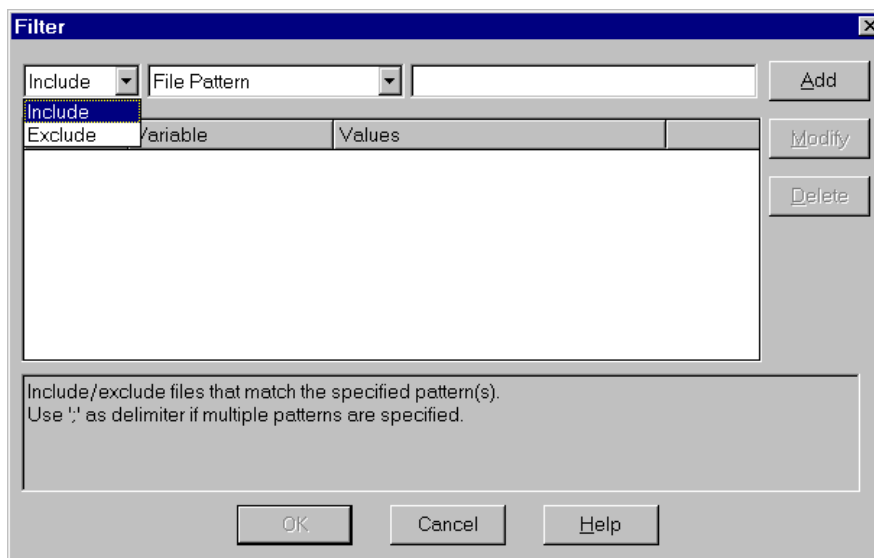
Backup Job Failure

If you determine that your job was not successful, you can use the information provided in the Last Result field to remedy the problem. You can then resubmit the job.

If a Job has a status of *Crashed*, it can be retried after the Job Engine restarts. This can be set up through the ARCserve’s Server Admin in the Job Engine Configuration tab. For further information, refer to the chapter “Administering the ARCserve Server” later in this guide.

Filters

Filters allow you to include or exclude files and directories from your backup and restore jobs, as well as from the utilities, such as copy, count and purge. You can access the filter options from the Backup Manager, Restore Manager, Copy, Count and Purge Utility windows.



Exclude

Exclusions always take precedence over inclusions. For example, if you add a filter to include files that have an .EXE extension, and you add another filter to exclude your \SYSTEM directory, all .EXE files in the \SYSTEM directory will be excluded.

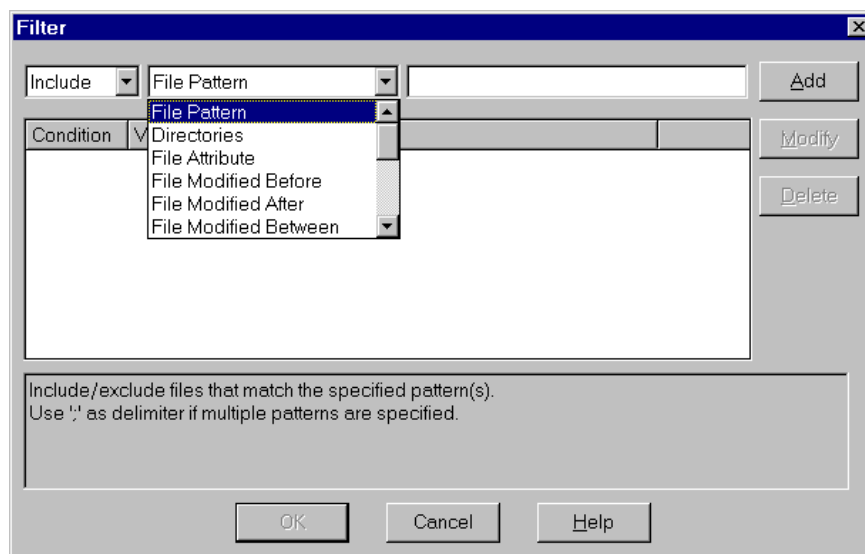
Include

Restore results contain only those files which satisfy the filter specifications. For example, suppose you selected to back up your entire local hard drive, and you then set up a filter to include files in the \SYSTEM directory. The result would be that ARCserve would only back up files from your \SYSTEM directory. No other files would be backed up.

Types of Filters

You have the option of including and excluding files based on the following criteria:

- Specific file names or patterns
- Specific directory names or patterns
- File attributes (hidden, system, etc.)
- Files modified before, after, between or within a specific date range
- Files created before, after, between or within a specific date range
- Files accessed before, after, between or within a specific date range
- File size equal to, greater than, less than, or between



For details on how to apply filters, refer to the online *ARCserve Procedures*.

File Pattern Filter

Use the File Pattern filter to include or exclude certain files from a job. You can specify a particular file name or you can use wildcards to specify a file pattern.

Directory Filter

Use the Directory filter to include or exclude specific directories from a job. You can enter an entire directory name or provide a pattern that the directory name follows.

File Attributes Filter

Use the File Attributes filter to include or exclude specific types of files from a job. There are five types of file attributes you can choose from. Select as many file types as you want.

- Hidden – Files not shown in a directory listing. i.e. IO.SYS is a hidden file.
- System – Files that are unique to the machine you are using.
- Archive – Files whose archive bit is set.
- Read Only – Files that cannot be modified.
- Offline – Files that are off-line.
- Migrated – Files that are migrated, if the Data Migration option is installed.

File Modified Filters

Use the files last modified attribute to include or exclude files based on the time they were last changed. There are four options from which to choose:

- Before – Files whose date matches, or whose date is earlier than, this date will be included or excluded.
- After – Files whose date matches, or whose date is later than, this date will be included or excluded.
- Between – Files whose date falls between the two dates will be included or excluded from the job. You must specify two dates for this selection.
- Within – Files whose date falls within the specified time will be included or excluded from the job. You must specify the number of days(s), month(s) or year(s).

File Created Filters

Use the files last created attribute to include or exclude files based on when they were created. There are four options from which to choose:

- Before – Files whose date matches, or whose date is earlier than, this date will be included or excluded.
- After – Files whose date matches, or whose date is later than, this date will be included or excluded.
- Between – Files whose date falls between the two dates will be included or excluded from the job. You must specify two dates for this selection.
- Within – Files whose date falls within the specified time will be included or excluded from the job. You must specify the number of days(s), month(s) or year(s).

File Accessed Filters

Use the file last accessed attribute to include or exclude files based on when they were last accessed. There are four options from which to choose:

- Before – Files whose date matches, or whose date is earlier than, this date will be included or excluded.
- After – Files whose date matches, or whose date is later than, this date will be included or excluded.
- Between – Files whose date falls between the two dates will be included or excluded from the job. You must specify two dates for this selection.
- Within – Files whose date falls within the specified time will be included or excluded from the job. You must specify the number of days(s), month(s) or year(s).

File Size Filters

Use the file size attribute to include or exclude files based on the specific size of the file. You can specify a size range from 0 to 9999999999 bytes, KB, MB or GB. There are four options from which to choose:

- Equal to – Files whose size matches will be included or excluded from the job.
- Greater than – Files whose size matches or whose size is greater than will be included or excluded from the job.
- Less than – Files whose size matches or whose size is less than will be included or excluded from the job.
- Between – Files whose size falls between the two sizes will be included or excluded from the job.

NDS Context & Object Filters

Use an NDS Context and Object filter to include or exclude certain NDS objects from your job.

Scheduling

All jobs can all be scheduled in essentially the same way, using the Schedule options available in each Manager. Backup jobs can also be configured to use Media Rotation schemes. Refer to the chapter, “Media Pools and Rotation Schemes” in this guide. Jobs can be submitted with a repeat method of:

- Once— Do not repeat this job.
- Every n days— Repeat this job every specified number of days.
- Day(s) of the Week— Repeat this job on the day(s) that are checked off.
- Week(s) of the Month— Repeat this job on the week(s) that are checked off.
- Day of the Month— Repeat this job on the specified day.
- Custom— Repeat this job on the month, day, hour or minute specified.

You should submit a *Run Job Now* if your storage device is currently busy. ARCserve will report that the storage device is busy and the backup job will not be submitted to the job queue. You will be asked to wait until the device is available before you can resubmit the job.

Instead, you should schedule your backup job, keeping the current date and time. This way, when ARCserve discovers that the storage device is busy, it will automatically retry the backup job until the drive becomes free.

You should submit a *Run Job Now* when:

- The job you are submitting is a “one-time” job that you want executed immediately.
- You want to monitor the job as it runs.

You should schedule your job when:

- You are submitting a “one-time” job, but you don’t want it to run now. You want it to run at a specific time.
- You are submitting a job that should run regularly. This is especially useful for setting up a media rotation scheme for your network.
- Your storage device is currently busy and you want to run a backup job as soon as the drive is free. To do this, schedule your backup job with the current date and time.

For details on how to specify a scheduling option, refer to ARCserve Procedures.

Job Scripts

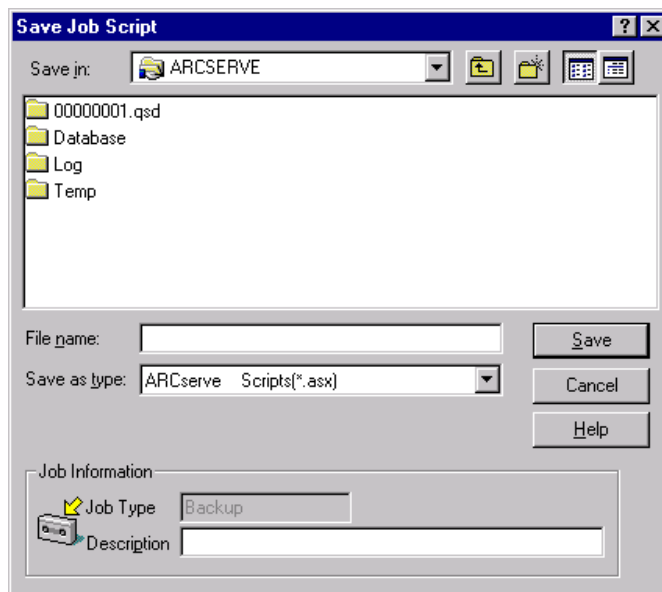
A script is a job that you saved to a file. It contains the original source, destination, options, and schedule information for the job. It will also contain any filters you created to include and exclude files and directories.

Creating a script has the following advantages:

- You can re-use the same settings at a later time.
- You can copy your settings to a different Windows 2000 machine running ARCserve.
- You can quickly re-submit regularly executed jobs after a job has been accidentally deleted.

Saving Scripts

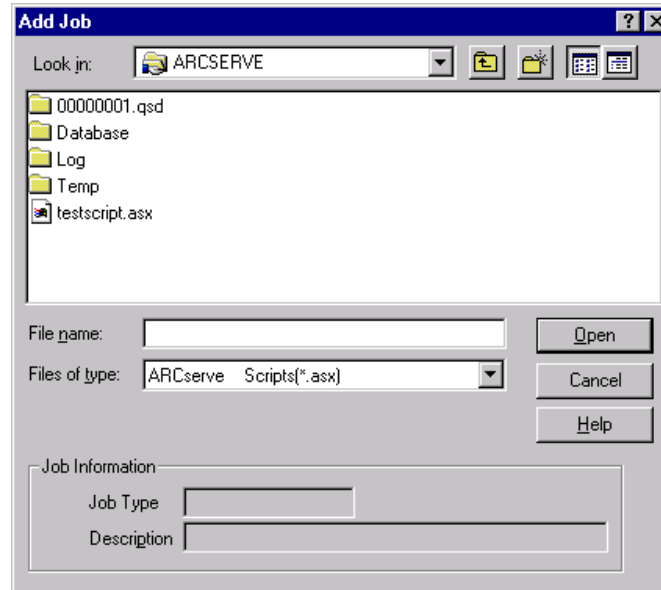
You can save any type of job as a script. After you create the job, choose the Run/Schedule button. In the Submit Job dialog, if you click the *Save* button, the following dialog allows you to save the job as a script:



After you save the script, select *OK* to submit the job to the queue.

Using an Existing Script

To use an existing script, go to the Job Queue tab of the Job Status Manager and click on the Add Job button. The following screen appears:



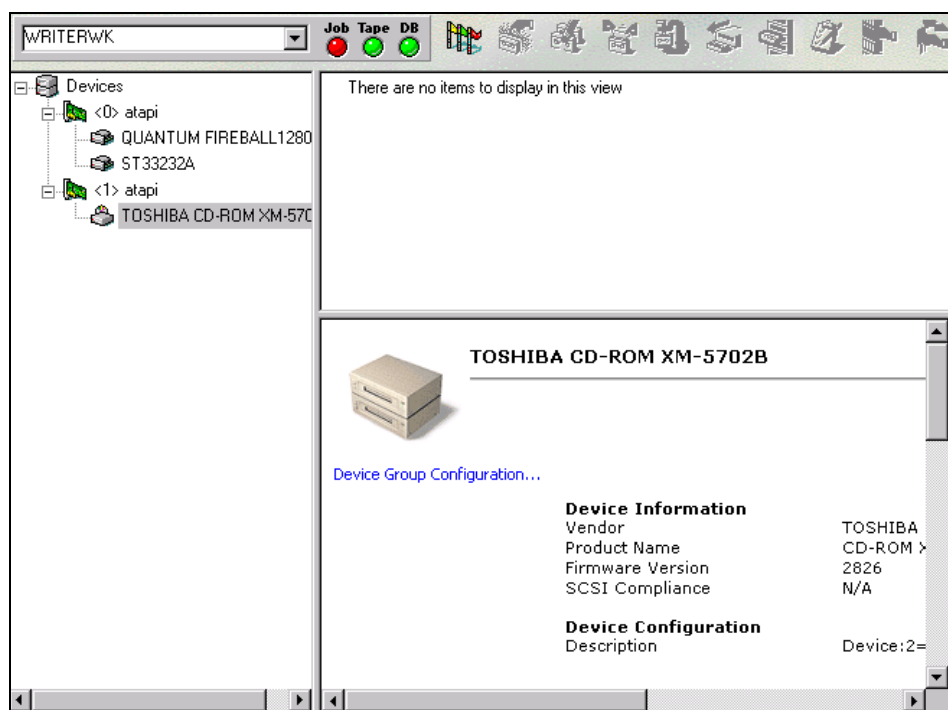
Choose the path and filename of an existing script, and then click on *Open*. The job information for the previously saved script will be displayed in the job queue as a new job.

For details on how to create and use a script, refer to the online *ARCserve Procedures*.

Managing Devices and Media

Device Manager Basics

The Device Manager provides you with information about storage devices that are connected to your system, the media in these devices, and the status of these devices.



The Device Manager also allows you to perform the following maintenance tasks on your media:

- Formatting
- Erasing data
- Retensioning

- Compressing data
- Ejecting
- Rebuild
- Inventory/Offline slots
- Mount/Dismount
- Import/Export
- Clean Tape heads

Important! Before you use these options, especially the destructive ones (formatting and erasing) make sure you have the right media selected.

The Device Manager is the starting point for all media and device monitoring and maintenance operations. Each device option is discussed in more detail in the following sections.

When you highlight a storage device, or the adapter card it is configured to, ARCserve displays summary information about the adapter card or the storage device, such as the vendor and model name, board configuration, etc.

Device Management Functions for Libraries

The Device Manager provides the following device management functions for libraries:

- Barcode/serial number media labeling
- Inventory Slot Range
- Format Slot Range
- Erase Slot Range
- Eject media from the library drives
- Mount Magazine
- Dismount Magazine
- Clean Library Drive (for tape libraries only)
- Import/Export
- Configuring Library Groups

Inserting and removing media

When you insert a media into a magazine slot or remove a media from a slot, you must either inventory this slot or remount the magazine.

Important: If you are manually inserting media into a library, always insert media into slots, *never* into its library drives.

Each of the Device Management functions is described in more detail below. For information on how to use the Device Manager, refer to the online *ARCserve Procedures*. For further information on the Device Management functions for libraries, refer to the Library Option Guide.

Formatting Media

You must format your media before using it. Although ARCserve automatically formats blank media, you can use this option to manually format your media. Formatting writes a new label at the beginning of the media, effectively destroying all existing data on the media. If you are using a library drive, you can select to format a single media or multiple media.

You can also use the Format option to assign media to media pools. For more information on this topic, refer to the chapter on media pools.

Important! Use this feature with care! Once you format media, the data it contained (if any) and any job sessions associated with this media are gone permanently.

A media's life is generally based on *passes*. A pass is defined as the storage device head "passing" over a given point on the media. For example, a backup without verification constitutes one pass, whereas a backup with verification constitutes two passes.

Media manufacturers rate their media's useful lives from about 500 to 1500 passes. This doesn't mean that the media is unusable after it reaches the maximum number, just that it is more susceptible to errors at this point.

The expiration date is a way of tracking how long the media has been in service so you can stop using it before it reaches the end of its useful life.

If you are formatting new, blank media, the default expiration date will be one year from the current date. If you are reformatting media, the expiration date that appears will be the date you specified the first time the media was formatted.

You should choose an expiration date based on how you will use the media. If you plan to use the media often, say, a few times a week, you should set the expiration date to a year from now, maybe even sooner than that. On the other hand, if you plan to use the media only once or twice a month, you can set the expiration date to two or three years from now.

When a media reaches its expiration date, you will still be able to use it, but when you use it to make a backup, for example, a note is made in the Activity Log that this media is expired.

Erasing Media

Use this option to erase all data from media. ARCserve also erases all references to the contents of this media (if any) from the database. If you re-format this media, its physical history (read and write passes, etc.) will be carried over. If you are using a library drive, you can select to erase data from a single media or from multiple media. There are three erase options from which you can choose: Quick Erase, Quick Erase Plus, and Long Erase.

***Important!** Use this feature with care! Once you erase data on media, the data is gone.*

Quick Erase

A Quick Erase takes less time than a Long Erase because it only overwrites the current media label. Although, technically, there is still data on the media, the data is effectively gone without the media label. Quick Erase is useful if you want to re-use ARCserve's media, but you don't have time to wait for a Long Erase.

Quick Erase Plus

Like Quick Erase, Quick Erase Plus erases the current media label. In addition, Quick Erase Plus also erases the serial number, so that you can assign a new serial number to the media.

Long Erase

A Long Erase completely removes all data from the media. It takes much longer than a Quick Erase, but the media is considered blank, as if it were just formatted. For security reasons, if you want to make sure that the data on the media is gone completely, use Long Erase. If you reformat removable storage media using the Long Erase with NTFS option, all existing data on the storage media will be erased.

Retensioning a Tape

Note: This feature applies only to Quarter Inch Cartridge tapes.

Use this feature to make sure a tape is evenly wound and properly tensioned to avoid errors, jamming, or worse yet, breaking. Retention a tape if you are having trouble writing to it or reading from it.

Compressing Data on Tape

Note: The storage device must support tape compression.

If your tape drive supports compression, you can tell ARCserve to turn it on. The only time you should turn it off is if you plan to use this tape in another drive that doesn't support compression, because the drive that doesn't support compression will not be able to read the compressed data on the tape.

Important! You can only change compression when a blank tape is in the drive. This prevents “mixing” of uncompressed and compressed data on a tape.

Ejecting Media

Use this function to eject media from a storage device. If you are using a library drive and the library is highlighted, it will eject the media in all of the drives. If you select a particular storage device, it will only eject from that drive.

Important! This task can only be used if the storage device supports this feature.

Rebuild

Note: This option applies only to the RAID Option.

This option rebuilds a lost or damaged media. ARCserve will rebuild the target media.

Inventory/Offline Slots

Note: This option applies only to libraries.

This option checks library slots to see if they are occupied, reads the media headers of any media it finds, then associates the media header of each media with the slot in which it was found. The slot in which each media was found is its home slot. You should use this option to inform the Tape Engine of any alterations made to the media in the library. The following options exist:

- Quick inventory – retrieve media information from the database.
- Regular inventory – read in all media information.
- Offline – dismount the media in the slot.

Note: You should only add and remove media when the Tape Engine is running so you can immediately inventory your slots.

Mount/Dismount

Note: This option applies only to libraries.

Use this option to load or remove a magazine from the library. Mounting a magazine initiates an inventory of the slots in the magazine. Dismounting a magazine returns all media to their home slots and prepares the magazine for removal. The time this process takes will vary according to the number of media in the magazine you are mounting or dismounting.

Note: Magazines must be mounted for library operations to commence. Magazines should be dismounted before you physically remove them.

Import/Export

Note: This option applies only to libraries.

The Import/Export Media screen allows you to add a new media to a large capacity library by specifying an empty slot to which the media can be assigned (import), or remove a media from a large capacity library (export). When you import media, the library reads the media and adds it to its inventory. Exporting makes it easy to remove a media for off-site storage, or to remove a media you suspect is bad. This screen contains the following fields and buttons:

- Mail Slot— Select an empty slot to import a media, or a slot which contains the media you want to export.
- Import— Click this button to import media.
- Export— Click this button to export the media.
- OK— Click this button when you are done.
- Quick import— Retrieve media information from database.
- Regular import— Read in all media information.

Clean Tape Heads

Note: This option applies only to tape libraries.

Clean the heads of any drive in your tape library. You must have a cleaning tape installed in the tape cleaning slot you specified during setup to utilize this option.

Device Groups

If you have more than one storage device connected to your machine, ARCserve lets you separate them into two or more *groups* (must be the same make and model). Establishing device groups is a key component to ARCserve's flexibility and efficiency. Device groups allow for:

- Parallel streaming - You can have several operations occurring simultaneously, one at each device configured for your system. For example, suppose you have two devices, DEV1 and DEV2, respectively. You could then have a backup job running with DEV1, while a restore job runs with DEV2.
- Automated media spanning - When one media becomes full, the session will automatically span to another media in another device within the same device group. Note that this requires two or more devices in the device group.

For example, suppose you have two device groups, GROUP1 and GROUP2. GROUP1 is made up of one storage device and GROUP2 is made up of two storage devices. If you submitted a backup job that will take up two media, you can insert blank media in each GROUP2 drive and let ARCserve automate the media spanning for you. If you used GROUP1 for the same backup job, you will need to manually eject the first media and then insert a second media. As you can see, automated media spanning allows ARCserve to do most of the work for you.

Note: By default, ARCserve is installed with each storage device assigned to its own group. If ARCserve finds identical storage devices (same make, model, and firmware), it automatically places them in the same group. Later, you can use the Device Manager to regroup your devices.

Configuring Device Groups

When ARCserve's Tape Engine is started, each storage device that you have is automatically placed in a device group. Identical drives are grouped together and different drives are separated into different groups. For example, if you have five drives attached to your machine, three of which are identical, you will have three groups: GROUP0, GROUP1, and GROUP2. The three identical drives will be in one group and the remaining two, since they are different from all others, are each placed in their own individual group.

Since ARCserve automatically groups your storage devices, you can keep these device group names or you can regroup and rename them. Use the Device Manager to:

- Create a new device group
- Assign a device to a device group (including a RAID group)
- Remove a storage device from a device group
- Rename or delete a device group
- Use a RAID tape set as one unit

For details on how to create new Device Groups, refer to ARCserve Procedures.

Removable Drive Support

ARCserve supports SCSI removable devices allowing you to back up data, restore data, scan sessions, merge removable sessions, and manage removable media on your removable devices. The ARCserve Manager identifies and treats the removable media as tape media.

Note: For the most up-to-date list of certified devices, go to Support from the ARCserve program.

Before you can back up to removable media, you must:

- Enable your removable device
- Format your removable storage media
- Configure removable device groups

Enabling Removable Drives

Before you can back up to your removable media, you must enable the drive so ARCserve can recognize it. This is done by clicking the Disk Configuration icon and configuring the removable drive by using the Administrative Tools.

Formatting Removable Media

After your device is recognized by ARCserve, you must format your removable storage media as an ARCserve storage media. In ARCserve's graphical user interface, removable media is represented as if it were tape media. This is not an error. ARCserve treats removable media in the same manner as tape media.

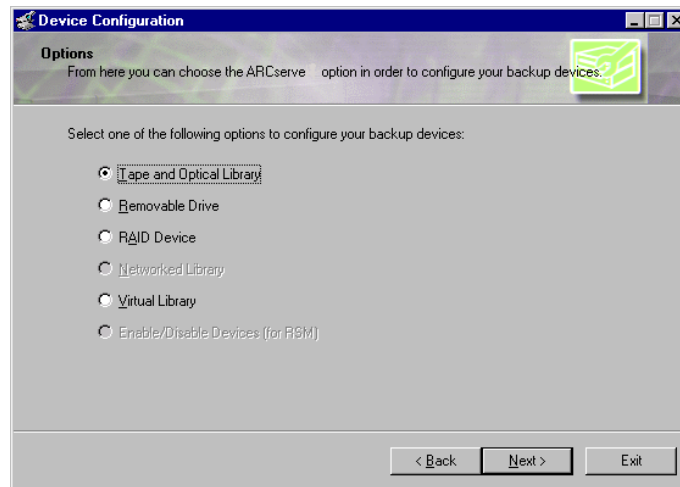
Configuring Removable Devices

You configure removable device groups through ARCserve's Device Management feature. Creating or deleting new removable device groups, renaming removable device groups, and assigning or removing individual devices are done easily with ARCserve. You cannot assign a removable drive into a group of tape drives. You must create a new group for the removable devices.

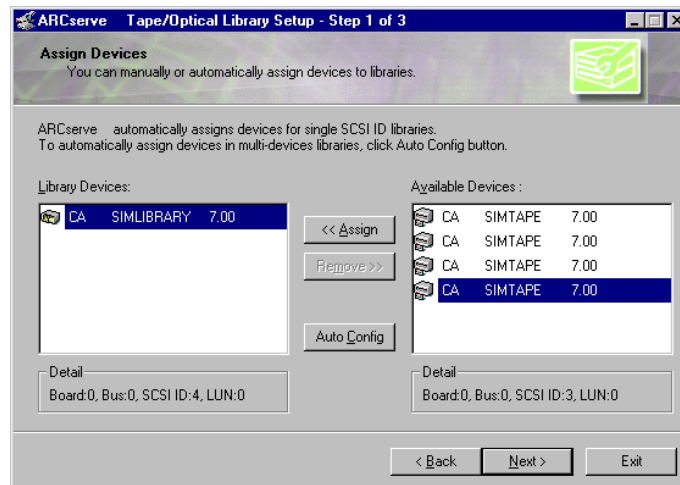
Configuring Libraries

You can configure (or modify the configuration of) a tape library by selecting Device Configuration in the ARCserve Wizards screen or in the ARCserve Program Group.

1. Choose Tape and Optical Library from the Device Configuration dialog box:



2. Click *Next*. The ARCserve Tape/Optical Library Setup screen should appear. **Note:** If your tape engine is started, a dialog appears, which will let you stop the engine. The engine must be stopped in order to configure a library device.



3. Assign the drives inside the libraries to the respective library in which it resides, using the Auto Config button, or by manually highlighting and assigning the drives.

Assigning a drive to a library allows ARCserve to be aware of the drive's existence within the library. If you are assigning drives manually, please note that drives inside multi-drive libraries must be assigned in the correct order.

Usually the library is configured by the manufacturer in such a way that the first library drive has the lowest SCSI ID number, and the last library drive has the highest SCSI ID number. **Note:** Be aware that this is not always the case – refer to the documentation supplied with your library for information on how your library's drives are configured.

4. Click *Next* - the next setup window opens. Highlight the library and click the *Properties* button. This allows you to specify which slot holds the cleaning tape (if present) and if a bar code reader is installed.

Note: This is for tape libraries only.

5. Click *OK*.
6. Click *Next*. The following setup window opens. This window provides the following advanced options:

- Library Quick Initialization – this allows ARCserve to skip the library inventory process that is automatically activated when the Tape Engine is started. This feature assumes that no media has been moved, added, removed, or swapped since the last shutdown. If any media has been moved, added, removed, or swapped, you should either manually inventory the entire library, or the specific slots whose status has changed. When the Quick Initialization feature is disabled, the entire library is inventoried each time ARCserve is started.

Note: ARCserve must run an inventory after configuring a library. Quick initialization will take effect after this first inventory is run.

- Mark Drive(s) Online/Offline – this allows you to mark defective drives (in a multi-drive library) as offline. ARCserve will continue to operate using the remaining non-defective drives. When the drive is repaired or replaced, it can be brought online.
7. Click *Next*. The setup-summary dialog box displays the choices you have made during setup.
 8. Click *Finish*.

Media Pools and Rotation Schemes

Media Pool Basics

A media pool is a collection of media managed as a unit. Each media pool is assigned a name and the media are organized according to serial numbers. The serial numbers assigned are permanent and if you are using a device with a bar code reader, the bar code labels will be used for the serial number of the media. The media pools are organized by the range of serial numbers of the media they contain and are divided into two sets, the Save Set and the Scratch Set.

The set of media containing important data that cannot be overwritten is called the Save Set. Once the media has passed certain criteria in a save set, they will be recycled and re-used.

Those media from the Save Set that can be re-used and overwritten are placed in the Scratch Set because they have met several criteria (minimum number of media in Save Set and retention period).

Each time a media in the Scratch Set is written to, it moves from the Scratch Set to the Save Set. That media will move back to the Scratch Set once the specified criteria have been met.

Media pools apply to every media, regardless of which backup type and method were selected. All Rotation backup jobs will create their own media pool based on the name entered in the Media Pool Name field of the Schedule dialog box for simple rotation, or the prefix entered in the Media Pool Name Prefix field with GFS enabled. The GFS backup uses three media pools: Daily, Weekly, and Monthly. These three pools are also based on the information entered in the Media Pool Name Prefix field. For example, if you enter ACTG as the Media Pool Name Prefix for a GFS backup job, the Daily media pool name would be ACTG_DLY, the Monthly name would be ACTG_MLY, and weekly ACTG_WLY.

For a simple (single media pool) rotation, you specify the complete name for the media pool.

Save Set

The media pool Save Set is a set of media which can not be overwritten. You can modify Save Set information for all Custom backup jobs. You can move media from the Save Set to the Scratch Set and you can move media from one media pool Save Set to another media pool Save Set.

You define the minimum number of media contained within the Save Set and the retention period. After the retention time for a media expires, it is eventually rotated to the Scratch Set in its pool.

ARCserve performs Media Pool maintenance at the beginning of a job, and will not allow media in the Save Set to be moved to the Scratch Set until two criteria are met:

- The oldest tape in the Save Set is compared to the retention time and exceeds it.
- The minimum required number of media is in the Save Set.

Otherwise, ARCserve prompts for a blank tape, or will accept media from the Scratch Set.

Note: In the Media Pool Manager you can view the media available in a Save Set or a Scratch Set by clicking on the icon for each type of set in a media pool.

Scratch Set

The media pool Scratch Set is a set of media which has been recycled from the Save Set after its retention period has passed. All media in the Scratch Set will be overwritten the next time they are used. Those media which have not been formatted in the longest period of time, will be used first.

Retention period

The retention period is the number of days in which a media has not been used before it is moved into the Scratch Set. For example, if you specify a retention period of 14 days, a media will remain in the Save Set as long as it has been used within that specified time. If the media has not been used for 14 days, it will be moved to the Scratch Set.

You can specify the minimum number of save media to be retained in the media pool. This is the number of media to be retained in the Save Set before the older media are recycled to the Scratch Set. This is a safeguard for preventing data loss in case backups are not done for extended periods of time.

Media Pool Manager

The Media Pool Manager allows you to create and maintain media pools, which allows a logical grouping of media to efficiently schedule the maintenance and recycling of your media. ARCserve creates a catalog file on your media to improve performance for merge jobs and database backup jobs. A serial number and/or a bar code can be assigned to a particular media to help manage your rotations. The Media Pool Manager helps you to perform the following tasks:

- Create a new media pool.
- Delete an existing media pool.
- Move media in a pool from Scratch Set to Save Set and vice versa.
- Perform location maintenance.
- Assign media to a media pool.
- Remove media from a media pool.
- Refresh view.
- View and reconfigure media pool properties such as:
 - Media name, serial number, and Random ID number
 - Name of media pool to which media belongs
 - Expiration date of media
 - Date first formatted
 - Date last formatted
 - Date media was sent
 - Date media was last written to
 - Date media was last read
 - Backup method
 - Set Status
 - Media Status
 - Location Status
 - Offsite Location

Creating a Media Pool

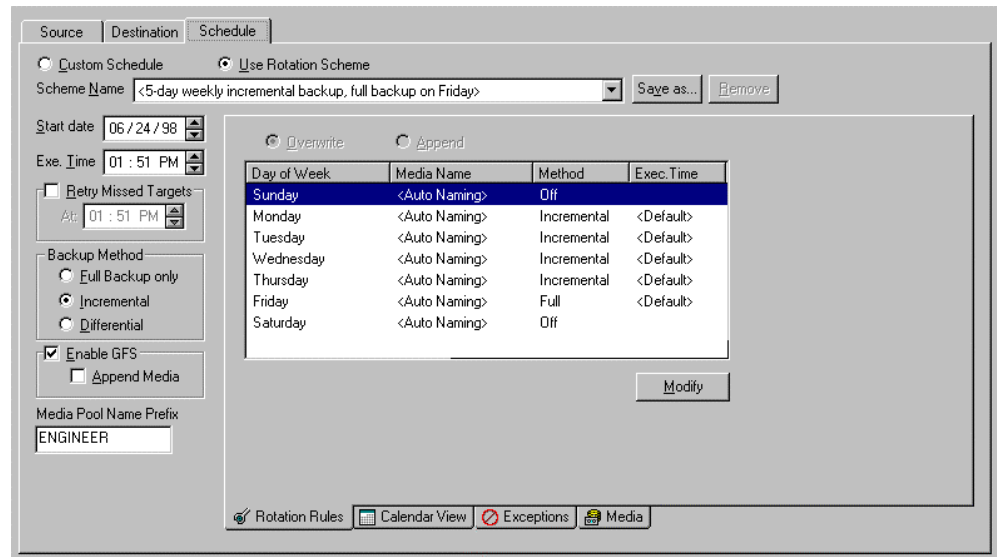
To assign media to a media pool, you first need to create the media pool. You can create media pools using the Media Pool Manager. A media pool name can consist of up to 16 uppercase characters. ARCserve creates a media pool(s), if they don't already exist when you submit a backup job either with simple or GFS rotation.

Tip: The Media menu in the Media Pool Manager dialog can be invoked as a pop-up menu by selecting an item and then clicking the right mouse button.

Deleting a Media Pool	You can delete a media pool using the Media Pool Manager. To delete a media pool, you must first re-assign the media to another media pool.
Moving Media	You can move media from one set to another. You can also move media from Scratch Set to Save Set and vice versa by using the Assign Media and Remove Media options from the context sensitive menu. Right-click on the appropriate set to view this menu. The media you've moved appear in their new locations.
Location Maintenance	You can assign or reassign offsite locations to media using the Media Manager. You can also enter new locations or modify the information about existing locations.
Assigning Media to a Media Pool	<p>You can assign media to a media pool during the process of formatting. When you format media using Device Management, you define certain media pooling information that will be associated with the media:</p> <ul style="list-style-type: none">■ New Media Name— If the media inserted in the storage device is named, that name will be displayed. If the media is blank, you can assign a media name. If the media is blank but has an assigned serial number, it can be assigned to any existing media pool.■ Expiration Date— You can set an expiration date for the media to be formatted. When this date arrives, ARCserve will remind you that the media has reached expiration.■ Media Pool— You can select any existing media pool previously created using the Media Pool Manager.■ Serial Number— The serial number is the current serial number, bar code label, or the next available serial number (based on pre-defined rules). You can manually enter a different serial number only if that number is not already in use. <p>Note: ARCserve scans the database, and if the serial number already exists (in any of the media pools), an error message appears prompting you to enter a new, unique serial number.</p>
Media Pool Properties	<p>When you highlight a Media Pool name and choose Properties, the dialog box that appears is the same as the dialog box that appears when you select New Pool except it contains information about an existing media pool. The fields are Pool Name, Base Serial Number, Next Serial Number, Serial Number range, Minimum # of Media in Save Set, Retention Time (days)</p> <p>When you highlight a media, you can view the following properties for any media in a media pool: Media name, serial number and Random ID number, Name of media pool to which media belongs, Expiration date of media, Date first formatted, Date last formatted, Date media was sent, Date media was last read</p>

Rotation and Grandfather-Father-Son (GFS) Media Pools

Rotation and Grandfather-Father-Son (GFS) media pools are based on basic media pooling architecture. The Rotation backup names its media pool based on what you enter in the Media Pool Name Prefix field of the Use Rotation Scheme view under the Schedule tab in the Backup Manager. For instance, the Media Pool Name Prefix in the following example is “Engineer”.



The GFS backup uses three media pools: Daily, Weekly, and Monthly. Each media pool has a default retention period and number of media to save.

Serial Number Information

The serial number of a media is one way to categorize media pools. You cannot change the serial number of media, but you can create a serial number for media by:

- Bar Code: A number is read from a bar code label, assigning the serial number. A changer with a bar code reader is required for this method. This will override any previously-defined media pool settings.
- Automatic: ARCserve automatically assigns a serial number for the media based on the Base and Range of serial numbers set when the pool was created.
 - Base: This is the base number which ARCserve will use when automatically assigning serial numbers. The first media formatted will have the same serial number as the Base number. Each media's serial number thereafter will be incremented by one.
 - Range: You can specify the range (up to 31 digits) from which the media pool serial numbers will be categorized.

Simple Rotation Scheme

This section describes how to configure a rotation scheme for a backup job by using ARCserve's default scheme or by specifying your own rotation parameters.

You can specify one of the following parameters for a backup on a given calendar day:

- Full backup - Performed each time the job is repeated.
- Incremental backup - Backs up only those files that have changed (their archive bit has been set) since the last Full or Incremental backup was done. After each backup, archive bits are reset so that they are not backed up during the next incremental backup job.
- Differential backup - The backup is performed daily, after one weekly full backup and the following occurs:
 - A full backup is made of the job and all archive bits are cleared.
 - The next day, the differential backup backs up files that have their archive bit set. (The archive bit will not be reset after the backup.)
 - The following day, the differential backup contains all files that have their archive bit set from the previous two days.
 - No backup

Note: You can use the Calendar View to customize your rotation scheme according to the types of backups you want for particular days of the week or month. This feature enables you to specify exceptions to the standard rotation scheme you're using.

GFS Rotation Schemes

The Grandfather-Father-Son (GFS) rotation scheme is a set of preconfigured backup jobs consisting of full backup jobs combined with incremental and differential jobs. ARCserve's GFS Rotation scheme is based on the standard backup schemes in the mainframe and minicomputer environments. This strategy allows you to base your backup jobs on a full weekly backup, then supplement that backup with daily incremental or differential backups.

Refer to the chapter "Backing Up Your Data" in this guide to review each backup method.

Incremental backups

On the Schedule tab of ARCserve's Backup Manager, you can choose a five or seven-day rotation scheme consisting of incremental or differential backups. In incremental backup jobs, the archive bit is cleared after every incremental backup job. Thus, only the files that have changed since the last incremental job are backed up again. In general, GFS rotation schemes that use incremental backup jobs require less time to perform the daily backup jobs during the week.

Differential backups

In GFS rotation schemes using differential daily backup jobs, only the files that have changed since the last full backup job are backed up. Since differential backup jobs do not clear a file's archive bit, the files that were backed up in the last differential job are backed up again. Using this scheme, the backup jobs require more time to process than incremental backup jobs. However, this strategy requires less effort to restore servers and workstations because you might need less tapes to restore your machines.

GFS Media Naming

When media is assigned to a GFS rotation media pool, ARCserve automatically assigns a name to the media, then labels the media with the backup type, date, and usage. This media naming convention allows administrators to easily identify backup media. Media for GFS rotation schemes are automatically named using the following syntax: (backup type)-(user-defined media pool prefix)-(day-of-the-week)-(date)

Where..	Is...
backup type	"F" - full backup "I" - incremental backup "D" - differential backup "W" - weekly backup "M" - monthly backup
user-defined media pool prefix	The name you assigned to the media pool for your GFS Rotation scheme.
day of the week	An abbreviation for the day of the week on which the job was performed.
Date	The date on which the backup was performed in mm/dd/yy format.

For example, the media used for the first backup in your rotation scheme will have the following name: F-TP-MON-11/1/00 where:

"F" denotes a full backup job

"TP" is the name of the media pool assigned for the GFS rotation scheme

"MON" stands for Monday, the day on which the rotation scheme starts

11/1/00 is the date on which the job runs

The GFS strategy is a method of maintaining backups on a daily, weekly, and monthly basis. The primary purpose of the GFS scheme is to suggest a minimum standard and consistent interval at which to rotate and retire the media.

GFS backup schemes are based on a five or seven-day weekly schedule beginning any day. A full backup is performed at least once a week. On all other days, full, partial, or no backups are performed. The daily backups are the Son. The last full backup in the week (the weekly backup) is the Father. The last full backup of the month (the monthly backup) is the Grandfather. GFS rotation schemes allow you to back up your servers for an entire year using a minimum of media.

Note: A five-day GFS rotation scheme requires 21 media-per-year, while a seven-day scheme requires 23 media-per-year.

Using GFS rotation, you can restore data reliably for any day of the week by using the weekly full backup in conjunction with the daily incremental or differential backup jobs.

You can deviate from your standard rotation scheme (for instance, if a holiday falls on Wednesday, your usual backup day). An exception is handled from the Backup Manager by selecting Use Rotation Scheme, and then selecting Exceptions. (Similarly, exceptions can be removed.)

GFS Media Pools

GFS rotation schemes are assigned to a particular group of media, which are recycled and reused throughout the cycle of the rotation scheme. Different media pools are created for your daily, weekly, and monthly backup media. ARCserve designs and maintains the individual media in the appropriate pool for the different backup types:

- Daily
- Weekly

- Monthly

For the media pool “BK” the following media pools are created in your GFS rotation schemes:

Pool name	Description
“BK_DLY”	Daily incremental, differential, or full backup media for the BK media pool.
“BK_WLY”	Weekly full backup media for the BK media pool.
“BK_MLY”	Monthly full backup media for the BK media pool.

Important! You must provide a Media Pool Name Prefix for your GFS rotation schemes.

By default, you can re-use daily media after six days. Weekly media can be overwritten after five weeks have passed since it was last written to. Monthly media is saved throughout the year. They can and should be taken off-site for storage. You can change any of these media rotation defaults to suit your particular environment.

Save and Scratch Sets

Save Sets and Scratch Sets are used for storing media in media pools. Media is assigned to Save Sets after a GFS backup job has been run and stays there until it is eligible to be assigned to a Scratch Set after the duration of the media’s retention time expires. Media in the scratch set is recycled reusing the oldest media in a scratch set first. Thus, in a GFS rotation scheme, a set number of media in a media pool is used for your backup jobs. If there is no media in a scratch set, ARCserve prompts for blank media.

The GFS rotation scheme automatically creates three media pools daily, weekly, and monthly, when you schedule a backup using this scheme. These pools are named according to the Media Pool Prefix you enter in the Schedule dialog box. Each media pool has its own Scratch Set and Save Set with a default retention period and minimum number of media to save.

In a GFS rotation scheme, each backup media is retained in the same set until the retention period expires. The media will be reused on the day after the retention period ends, provided that the minimum number of media exists in the save set.

The retention time for media depends on the type of GFS rotation scheme you are using, that is a five-day GFS rotation scheme or a seven-day rotation scheme, and the media pool to which the media is assigned, daily (_DLY), weekly (_WLY), or monthly (_MLY).

Five-day GFS rotation schemes have the following retention times for each media pool:

- Daily (_DLY) - six days (daily media in seven-day rotation schemes has a retention time of eight days)
- Weekly (_WLY) - five weeks
- Monthly (_MLY) - 343 days

If, for example, you have a daily incremental backup job for a five-day GFS rotation scheme on the media "I-SAMPLE-TUES-9/09/00", the media remains in the save set for six days. The media will be overwritten when the same job runs again the following Tuesday. The new name for the media will be "I-SAMPLE-TUES-9/14/00".

The following are the formulas used for calculating the number of media in the save sets and the retention times for the GFS media pools:

- Daily pool: This pool holds the media for daily backup jobs. The default retention period is six days and the number of Save Set media is based on the number of daily media in the GFS rotation minus one [# of daily media -1]. For example, if you are using the five-day GFS rotation scheme, in which you preserve six daily media, the number of Save Set media in the GFS Daily Media Pool would be five: six daily media minus one. After six days have passed since the media was used (on the seventh day), it meets the criteria to be moved from the Save Set to the Scratch Set.
- Weekly pool: This pool holds the weekly media. The retention period equals the number of weekly media times seven, minus one [(# of weeklies * 7) - 1]. The number of save media is based on the number of weekly media in the GFS setup minus one [# of weekly media - 1].

For example, if you are using the seven-day GFS rotation scheme, in which you preserve five weekly media, 34 days would be the length of the retention period (5 times 7 minus 1). In the same example, four would be the number of save media in the GFS Media Pool (5 weekly media minus 1). After 34 days since the media was used (on the 35th day), it can be moved to the Scratch Set, provided there are four other media in the Save Set of the Weekly pool.








- Monthly pool: This pool holds the monthly media. The retention period equals the number of monthly media times 29 minus five [(# of monthlies * 29) - 5]. The number of save media is based on the number of monthly media in the GFS setup minus one [# of monthly media - 1].

For example, if you are using the seven-day GFS rotation scheme, in which you preserve 12 monthly media, 343 days would be the length of the retention period (12 times 29 minus 5), which takes into account months with differing number of days. In the same example, 11 would be the number of save media in the GFS Media Pool (12 monthly media minus 1). After 343 days since a media has been used, it can be moved to the Scratch Set, provided there are 11 media in the Save Set of the Monthly pool.

Selecting GFS Rotation Schemes

ARCserve allows you to select from pre-defined GFS rotation schemes on the Schedule tab of the Backup Manager when you click the Use Rotation Scheme option. Although GFS rotation schemes are predefined, you can modify these schemes to suit your individual needs. You can use the following tabs to customize your GFS rotation jobs:

- Rotation Rules - Customize rotation rules. As an alternative to selecting Remove Exceptions from the Calendar View, you can remove exceptions using the Exceptions dialog.
- Calendar View - Customize individual days. With GFS rotation either enabled or disabled, you can use the Calendar View feature to customize your rotation scheme according to the types of backups you want for particular days of the week or month:

Sun	Mon	Tue	Wed	Thu	Fri	Sat
	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	 Full	 Incremental	 Full	 Off
 Off	 Incremental	 Incremental				

- Exceptions - Define particular days on which the back up type and the execution time or date differs from the pre-existing schemes. You can also select specific days off or define different backup criteria for days in your rotation on which you want to deviate from the standard scheme.

- Media - Append the new sessions to existing media with data. At the beginning of the week or for a full backup, ARCserve formats an appropriate tape (from the Scratch Set or a blank tape). For every Differential or Incremental backup *after* a full backup, ARCserve formats a blank tape or one from the Scratch Set.

Sample GFS Rotation

If you were using the five-day incremental GFS rotation scheme as your backup strategy, the media would be used in the following manner during one month:

Sun	Mon	Tue	Wed	Thu	Fri	Sat
	1	2	3 Full F-SAMPLE-WED... <Default>	4 Incremental I-SAMPLE-THU... <Default>	5 Full W-SAMPLE-FRI... <Default>	6 Off
7 Off	8 Incremental I-SAMPLE-MON... <Default>	9 Incremental I-SAMPLE-TUE... <Default>	10 Incremental I-SAMPLE-WED... <Default>	11 Incremental I-SAMPLE-THU... <Default>	12 Full W-SAMPLE-FRI... <Default>	13 Off
14 Off	15 Incremental I-SAMPLE-MON... <Default>	16 Incremental I-SAMPLE-TUE... <Default>	17 Incremental I-SAMPLE-WED... <Default>	18 Incremental I-SAMPLE-THU... <Default>	19 Full W-SAMPLE-FRI... <Default>	20 Off
21 Off	22 Incremental I-SAMPLE-MON... <Default>	23 Incremental I-SAMPLE-TUE... <Default>	24 Incremental I-SAMPLE-WED... <Default>	25 Incremental I-SAMPLE-THU... <Default>	26 Full M-SAMPLE-FRI... <Default>	27 Off
28 Off	29 Incremental I-SAMPLE-MON... <Default>	30 Incremental I-SAMPLE-TUE... <Default>				

Using RAID with ARCserve

RAID (Redundant Array of Inexpensive Devices) enables you to use multiple drives to process (back up, restore, etc.) data. Data is written across all the drives (or disks) used in the array in a predefined order. The drive array you set up will appear (on all informational screens) as a single drive. Internally, the multiple drives are accessed in parallel.

An ARCserve RAID is a set of drives that behave as a single, high-performance drive. Using ARCserve with RAID is similar to regular ARCserve operation, but with RAID you are dealing with a RAID tape set as a unit, instead of a single tape drive.

Notes:

1. RAID applies only to an array of tape drives (not hard disks).
2. RAID supports tape libraries.

RAID Levels

RAID makes the multiple drives you set up in your array appear as one “virtual” drive. In this way, RAID provides several data management advantages. RAID ensures fault tolerance and increases throughput. The following section describes and compares the various levels used by RAID:

RAID 1- Mirroring

Mirroring provides fault tolerance and reliability. In the mirroring mode the operations of the drives are duplicated. Whenever data is written to one drive, the same data is also written to the redundant drive, so there are always two copies of the data. This redundant drive ensures that the backup can be completed, even if a drive or media fails. In case of drive failure, the backup or restore continues with the remaining drive.

The tape sets used in mirroring are compatible with ARCserve (in non-RAID mode) functions. You can use ARCserve (non-RAID mode) to restore data from any of the RAID tapes, but it is not recommended to back up to any of the RAID tapes with ARCserve (in non-RAID mode).

NOTE: A session cannot be appended to a RAID-1 tape in a standalone drive.

RAID 0- Data Striping Data striping transparently distributes data over multiple drives to make them appear as a single, high-performance drive. It improves throughput performance by allowing data to be processed in a parallel manner. Although you have high performance with data striping, you do not have fault tolerance.

RAID 0- Striping (Not fault tolerant) Striping creates high throughput, but does not provide fault tolerance. RAID Level 0 is an ideal solution when you require optimum throughput rather than fault tolerance.

RAID 5- Parity Parity is an error-detection method, which is used to detect and correct errors in data transmissions. It is used to perform specific checks of the data blocks within a stripe. In the event of a single failure, the information for the missing drive or tape can be re-created by examining the remaining data and parity and regenerating the missing data bits.

With a RAID array using three or more drives, data and parity are striped across all the drives. Data is processed on two drives while the third drive holds the parity. If one of the media is lost, parity ensures that the missing data can be re-created.

RAID Fault Tolerance

The most important benefit of RAID is fault tolerance. During a restore, if one of the tapes becomes unusable, you can still restore from the remaining tape(s). RAID offers two main types of fault tolerance modes from which you can select:

RAID 1- Fault Tolerance Tape Mirroring This level provides fault tolerance and tape mirroring. The mirroring mode maintains both tape drives as exact duplicates. Tape mirroring ensures that the data will continue to be delivered to a drive even if the other drive fails. This RAID level is preferred if you want to maintain a tape at one site while sending the mirrored tape to a different location. You can transport your data easily without performing a tape copy operation.

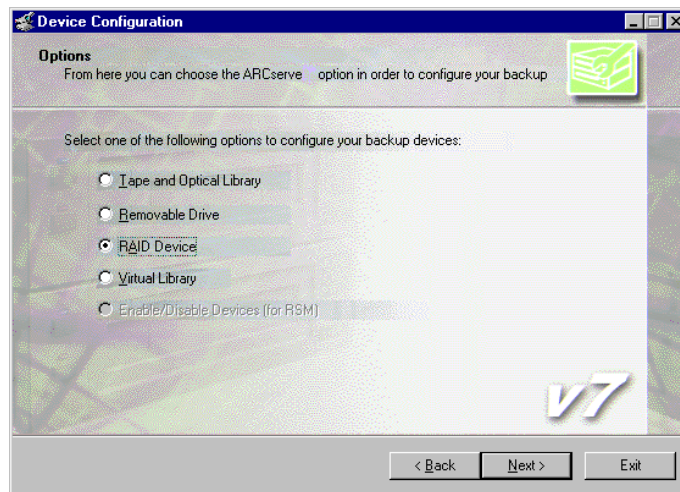
Tip: We recommend when using mirroring, you maintain one tape copy offsite for safekeeping.

RAID 5- Fault Tolerant Striping Parity Generation This level involves data striping. It improves reliability so that in case of failure of one drive, the parity (error correction) can be used to generate the lost data. Performance is improved because data is striped across three or more drives and parity (error correction information) is generated and stored.

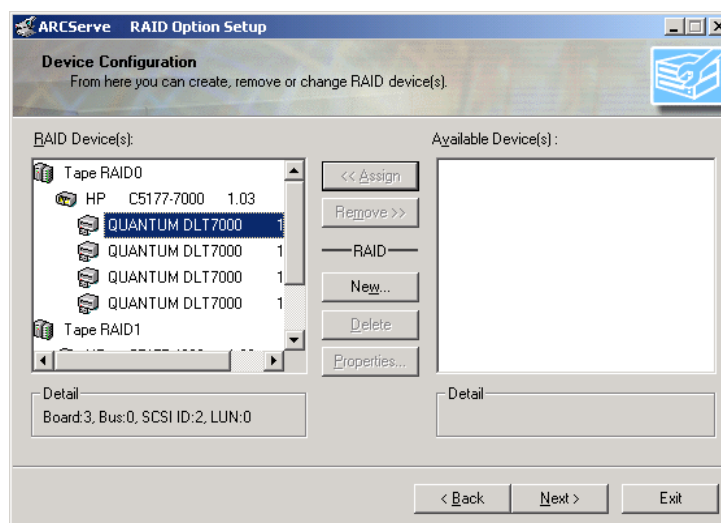
RAID 5 increases the “virtual” capacity of your tape (with RAID, three drives appear as one “virtual” drive). For example, if your tape’s capacity is two gigabytes, normally you would need two tapes (spanning) to back up four gigabytes. However, with RAID 5 you can store the same four gigabytes on one “virtual” tape (no spanning). If you are using three drives with RAID5 your storage capacity will double. If you are using five drives with RAID 5 your “perceived” storage capacity will quadruple.

Configuring RAID Levels

Configure a RAID level by selecting Device Configuration in the ARCserve program group. Choose RAID Device from the Device Configuration dialog box:



When you click Next, the RAID Option Setup dialog appears, enabling you to create a new RAID device, assign a RAID Level, delete an existing RAID, or change the RAID level:



We recommend that you review the attributes of each RAID level discussed earlier in this chapter, and refer to the ARCserve Procedures for online instruction on selecting a RAID level and assigning drives to the RAID device.

Device Configuration

The RAID device must be added to a group in the ARCserve Device Manager, to perform backup, restore, and copy operations using that RAID device. When running the job, ARCserve automatically assigns a RAID device to a group, if it is not already assigned. Refer to the online *ARCserve Procedures* for instructions on how to manually assign a RAID device to a RAID group.

You can use the ARCserve Device Manager window to format, erase, eject, or change compression for a RAID tape set. Refer to the online ARCserve Procedures for information to carry out these operations.

Important! All tapes in a RAID tape set are subject to each operation (formatting, erasing, ejecting, changing compression). For example, if you select Format, all tapes in the drives comprising the specified RAID device will be formatted concurrently.

Backing Up and Restoring with RAID

RAID's fault tolerance allows you to continue your backup or restore even if the media goes bad or if one drive fails. Using this feature improves backup and restore performance by distributing data over multiple drives.

The main difference using RAID, is that you are now dealing with a RAID tape set as a unit, instead of a single tape drive. RAID is displayed in ARCserve as a virtual drive, and the special processing of the RAID backup is transparent.

To back up data to a RAID device, select the RAID grouping as the destination in the Backup window and follow the standard backup procedures. To restore data from a RAID device, select the RAID grouping as the source in the Restore window and follow the standard restore procedures.

Maximizing Performance with RAID

ARCserve RAID offers significant performance benefits to your backup system by providing fault tolerance. You can monitor your system performance using the Windows 2000 Performance Monitor.

In order to maximize your performance, make sure of the following:

- Your data bus can handle the increased data flow.
- Your drive array can support the increased data flow.

Hardware and Software Requirements

Make sure that your hardware and software requirements match those recommended by Computer Associates.

- Use PCI/VLB or EISA adapters (i.e., Adaptec AHA-2940).
- Connect your tape drives to dedicated controllers.
- When using high performance drives, you may need to limit the number of drives per adapter.
- Use proper media type (DLT IV vs. DLT III, DDS3 vs. DDS).
- Make sure that your drives are identical.

Poor Performance

For some types of drives, connecting more than one drive to the same board can significantly reduce performance. There are other factors that could affect performance:

- The hard disk could be the bottleneck. Make sure you have a hard disk that is fast enough to feed data to the RAID.
- One of the tape drives is slow due to bad media.
- If another application accesses the same hard disk during backup, the performance will not be optimized because of hardware limitations.


Buffer Size

The performance of a RAID device will depend on the speed of the slowest drive attached. For example, in a 3-drive Level 0 RAID where drive 1=100 MB/min., drive 2=120 MB/min., and drive 3=120 MB/min., the total throughput will be 300 MB/min. Keep this information in mind when coordinating drives for RAID.

By increasing the buffer size, the performance of a backup or restore job can be improved. The optimum buffer size to set depends on a number of factors:

- the tape drive, the hard disk, and the amount of memory in the system
- the number of jobs that are running at once

You can modify the buffer size in the Create RAID Device dialog. If *Default* is specified, ARCserve will calculate an optimum value and use it. If *Custom* is specified, you can specify the buffer size. You can find out the best buffer size for your environment by trying a few different values.



Troubleshooting

Error incidents are listed in the Activity log and will have the prefix E with a number following it:

Problems	Possible Cause and Remedies
E6623: Level has been changed from Level 5 to Level 0.	When the ARCserve Tape engine initializes, it checks each RAID and makes sure that all RAID devices are configured correctly. If it finds a RAID that is configured as Level 5 and it doesn't find a minimum of 3 drives required for level 5 in RAID, it automatically changes the level to 0. (This can happen if by mistake one of the tape drives is not switched on when NT boots up. To access the devices switched on after NT boots up, run any utility that rescans the SCSI bus, for example WINDISK.EXE.
E6608: "Tape length mismatch."	One of the tape drives in the Level 1 RAID had failed during the last backup to the tape set, and, as a result, the data on one tape is less than the data on the other. ARCserve uses the tape with maximum data on it and continues the job.

Managing the Database and Reporting

The ARCserve Database records job information for each job that is run, including media and media device information. The database information allows you to perform restores by keeping track of each file and directory that was backed up to media. When you want to restore a specific file, the database determines which media a file is stored on.

The database information is also used to generate ARCserve's many types of reports. Consult your Computer Associates reseller for specific information about the databases that are supported.

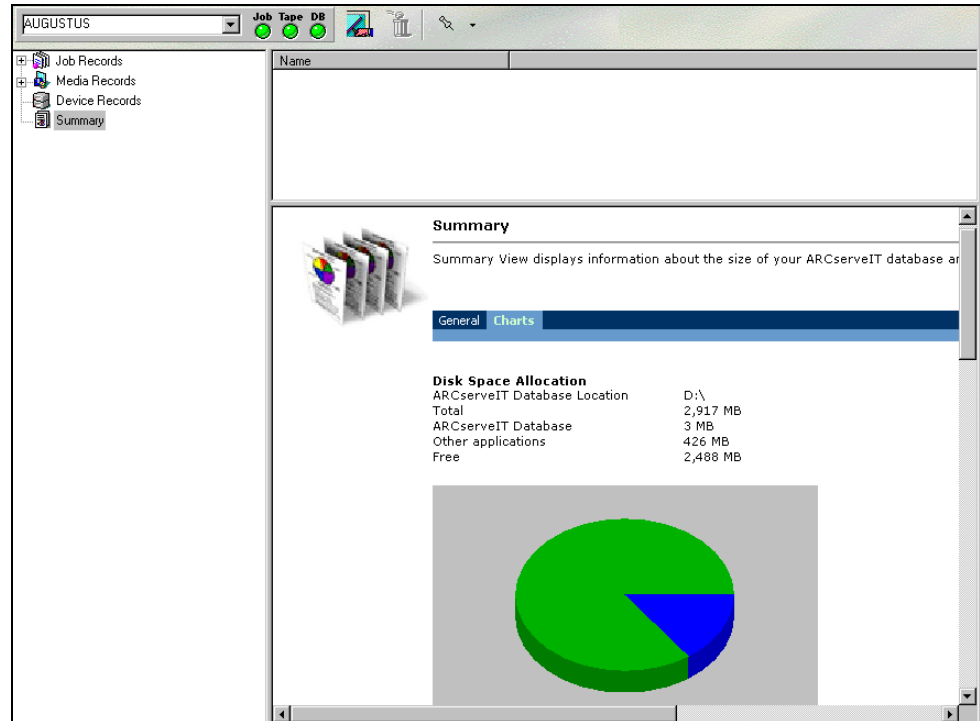
The Database Manager

The Database Manager allows you to:

- Keep track of the location of your media.
- Determine the session number of your backup.
- Determine if media should be retired.
- View log information about jobs you have run.
- Delete old records from the database.
- Graphically compare the size of your database to the total available disk space.

Note for SQL: The database size reported by the ARCserve Database Manager is the size of the data device (not including the log). You can obtain this information by entering the following command in the SQL Enterprise Manager: `sp_spaceused`

When you open the Database Manager from the Quick Access screen, it displays the current job records.



Selecting a Database View

When you open the Database Manager, the Left Panel displays four options:

- Job Records – displays information about all jobs that have been processed by ARCserve
- Media Records – displays information about all the media that have been used by ARCserve
- Device Records – displays information about all the devices that have been used by ARCserve
- Summary – displays how much space the database uses on your hard disk, database type and other settings.

Changing the Sort Order

You can change the sort order of records displayed in the Job Records, Media Records, and Device Records view by clicking on the field name you want to sort.

Pruning Your Database

You may want to set ARCserve to remove old records from the database. Refer to the chapter, “Administering the ARCserve Server” in this guide for more information.

Re-building the SQL Index

Note: This only applies when using a Microsoft SQL as the ARCserve database.

We recommend that you rebuild the index periodically to keep the index size manageable. The index should be updated after running an overnight job in order to keep the database size small. If you do not have enough time to update all the indexes, update the key index “IX_astpdat_1”. This index plays an important role and affects the browsing speed dramatically in the Restore Manager and Database Manager. Refer to the Microsoft SQL Online Help for steps on rebuilding the SQL Index. You should also update the index statistics from time to time.

Errors

The following describes the statistical information that is available about the database:

- Media Errors— Indicates that some sort of data corruption has occurred on the media and the read or write operation could not be successfully completed.
- Soft Read Errors— An error has occurred while reading the media. ARCserve has attempted to correct the problem in real time. A higher number of soft read errors indicates possible defective media. Media should be replaced for any future backups.
- Soft Write Errors— A write has occurred during the backup. ARCserve is correcting the media problem in real time. A high number indicates the media should be replaced for future backups. Make sure the drive heads are cleaned after the current backup session has been completed.

Device Records

If a drive has a critical error, the Error Log will contain some of the following information:

- Time— Time the error occurred.
- Sense Info— SCSI error code.
- Media— Number of media errors that occurred during the job.
- Soft Write— Number of soft write errors. (Refer to the explanation about the Statistics tab.)
- Soft Read— Number of soft read errors that occurred during the job. (Refer to the explanation about the Statistics tab.)
- Media Usage— Amount of time the media was used during the job.
- KB written— Amount of data written to the media during the job.
- Times Formatted— Number of times media has been formatted.

Using a Centralized Database

You can use ARCserve's central database to restore data, and use media and session detail for generating reports for central database members. The members, which are subordinate ARCserve machines on your network, update the central ARCserve database with media session information and details from their own databases. The central database is set up locally on a machine which manages the centralized database. Media session information for all of the ARCserve machines in your enterprise is contained there.

You can configure the relationships between the central database and the databases of the domain's other ARCserve machines in ARCserve's Central Database Administrator.

Tip: To increase the performance of your central database, we recommend installing it locally on your central ARCserve machine.

You can set up a Microsoft SQL database or ARCserve's Raima database as your central database. Administrators can use the central database to perform the following functions:

- Query for historical information from multiple backup machines from a central location.
- Generate adhoc reports for all of the individual ARCserve databases in your enterprise since all of your data resides in a central location.
- Restore machines from a central location in your enterprise.

The first step in setting up a centralized scheme for ARCserve on your network is to designate your central databases, as detailed in the following section.

Adding Members to an ARCserve Central Database

You must have Administrator rights to add members to a central database in order to establish an informational relationship between the parent database and the subordinate databases. The data flows from the dependent ARCserve databases to the ARCserve's central database. Dependent ARCserve machines can only send data to one central database.

If you select to write the job information to the ARCserve central database, the server that the central database resides on must be running.

Configuring the Job Information Stored in Your Central Database

In the Central Database Administrator, you can configure what job information to store locally or in the central database. The following table details the different database methods you can configure for the central and subordinate databases:

Method	Central Database	Local Database
1	Job summary information, including media session information, is recorded, but no job detail is stored.	No job details are recorded.
2	All job details are recorded in central database only.	No job details are recorded.
3	Job summary information, including media session information, is recorded, but no job detail is stored.	Summary and detailed job information is recorded on the databases of the subordinate ARCserve machines.
4	Summary and detailed job information is recorded in the central ARCserve database.	Summary and detailed job information is recorded on the databases of the subordinate ARCserve machines.

Using a Microsoft SQL Server

There are two database options when you install ARCserve: the standard ARCserve database (RAIMA/VLDB) or Microsoft SQL Server (6.5/7.0)

Note: For the SQL database to function properly, it must be started using Microsoft's SQL Service Manager. You also cannot use the restore by tree method if using the Microsoft SQL database.

Transferring Data from a Raima Database to a SQL Database

When configuring a Microsoft SQL database for ARCserve, you might want to use an existing SQL database for storing an ARCserve media session, job summary, and detailed job information. You can use the transfer utility (SQLIN.EXE) located in the ARCserve home directory to transfer this job information from an ARCserve standard database to a Microsoft SQL database. However, any information existing in the database is overwritten.

Important! You should not attempt to transfer information from a Raima database to an SQL database containing information, because the SQL database might be overwritten. Instead, you should create a new SQL database for storing your ARCserve media and job information.

SQL Connections

For each job that you run, you need two SQL connections. Be sure that you have set enough connections (or licenses) in your SQL server. To determine your default SQL connections, select Server and SQL server from the SQL Enterprise Manager. When you browse from the configuration tab, you will view the user connections. Set these values to the appropriate user setting. If an error message appears, “Cannot Update Record”, or “Failed to Login”, you may have run out of connections.

Note: Increase the open object to 2000.

Run a Database Consistency Check

When your database activity is not high, we recommend that you run a Database Consistency Check if you have a large database. Although it takes some time, it is important to determine that your SQL database is functioning well. For more information, refer to your Microsoft SQL guide.

The Report Manager

ARCserve provides the following options for displaying logs and reports:

- Activity Log: Logs all ARCserve activity.
- Tape Log: Logs all tape activity (for debugging purposes only)
- Job Log: Logs activity related to a specific job.
- Reports Manager: Generates reports from the ARCserve database for viewing or printing.

Important Note About Logs: Be sure to monitor the log size periodically to be sure that the log is not full. If a log is full, the database cannot function. Although the default setting is “truncate log on checkpoint”, you should increase the log size to 50% of the database if you expect to keep a large number of records.

The Activity Log

The Activity Log contains comprehensive information about the operations performed by ARCserve. It provides an audit trail of all ARCserve activity, including every job that is run. For each job, the log includes the following:

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

You can scan this log every day to see if any errors have occurred. You can also use it to find out a session number in case you need to restore a specific session.

The Tape Log

The Tape Log contains messages sent by the tape drives to ARCserve. This log is not generated during normal operation. It is designed for debugging purposes only.

The Job Log

A Job Log is generated for each job that is run by ARCserve. You can specify the level of detail in the log by choosing the Log options before you submit the job. The online *ARCserve Procedures* shows you how to configure and view the log report for a job. The log options are:

- Log All Activity – Record all of the activity that occurs while the job is running.
- Log Summary Only (default) – Record summary information on the job (including source, destination, session number, and totals) and errors.
- Log Disabled – Do not record any information about this job.

Viewing Reports

You can open the Reports Manager from the Quick Access menu. The Report Manager allows you to view and print a variety of reports based on information in the ARCserve database. The types of reports are:

- Job Report— Lists all of the jobs that have been run by ARCserve. For each job, this report lists Job ID, Submitter, Status, Start Date, Type and description.
- Backup Media Error Report— Provides information about the medium used with ARCserve. Information is reported when media errors are generated. This report lists the date the media was formatted, how often it has been used, the amount of data written to the media, and the number of errors that have occurred during use.
- Session Report— Provides information about all of the sessions that have been backed up on media. The report contains the session information that is found in the Database Manager on the Job Records or Media Records View when you highlight a record and click the Session tab. When you go to view this report you will be asked to select a media and a session from that media.
- Session Detail Report— Provides information about all of the files that were backed up in a particular session. The report contains the session information that is found in the Database Manager on the Job Records or Media Records View when you highlight a record and click the Files tab. When you go to view this report you will be asked to select a media and a session from that media.
- Backup Device Report— Provides information about the backup devices you are using with ARCserve. This report lists how much use the device has had, and the number of errors that have occurred.
- Enterprise Job Status Report— Provides information about the jobs that have been run by all the servers within a specific period of time. The information includes Server Name, Job ID, Start Time, Time Elapsed, Job Status and Submitter, along with the total number of files, total number of bytes, transfer rate of each server and the date of the session for each job.
- Media Pool Report— Provides information about the defined media pools. This information includes Pool Name, date activated, serial number ranges, as well as the media properties of the tapes contained in the pool.
- GFS Media Prediction Report— Provides the prediction (displays a forecast) of the media schedule for the selected GFS job. You can select the number of days you want to run the report and from the time at which you run it.

Administering the ARCserve Server

ARCserve Engines

The ARCserve Server consists of three functional engines or services:

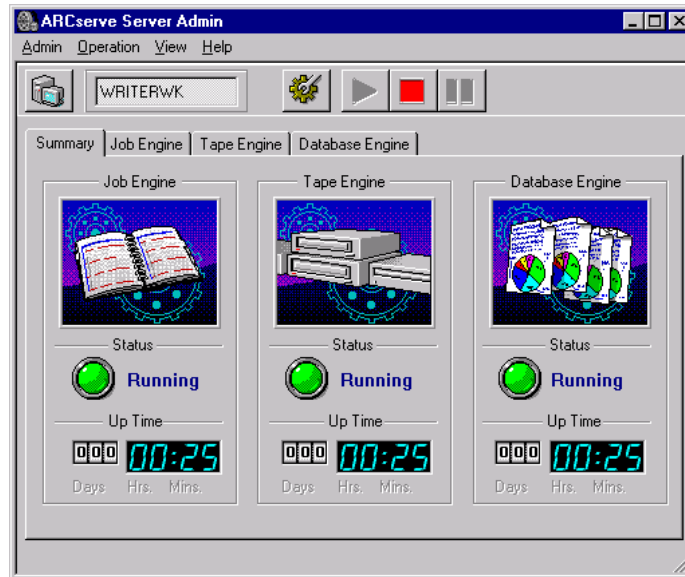
- The Job Engine – Processes your jobs at their designated date and time. It scans the job queue for a job that is ready to run, then sends it to the appropriate handler.
- The Tape Engine – Communicates with and controls your storage devices. The Tape Engine selects the device needed for an ARCserve job.
- The Database Engine – Maintains a history of:
 - Files, directories, drives, and machines that ARCserve has backed up or copied.
 - Information about jobs that have been processed by ARCserve, such as the job type, logs, the final result, and the start and end time.
 - Media used by ARCserve, such as its type, its name, the date it was first formatted, the date it expires, and the sessions on it. You can manage an ARCserve Server from a remote machine.

You can control these ARCserve services in the Server Admin. To view information about an individual engine, launch the ARCserve Server Admin program and select the particular engine menu tab.

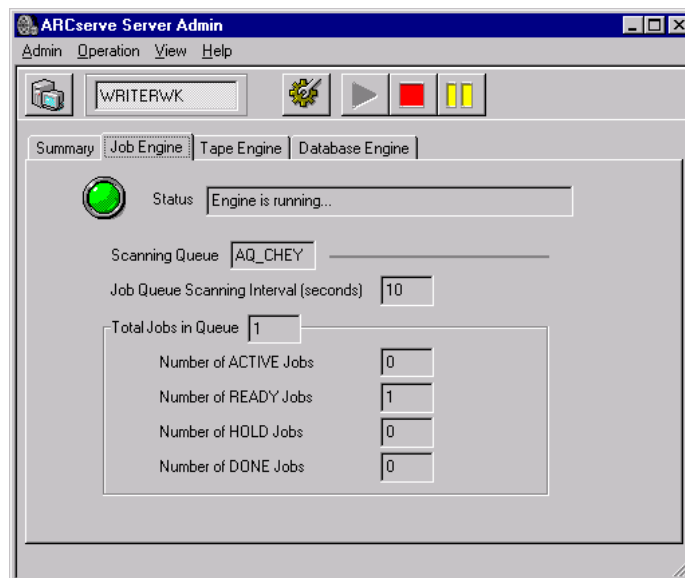
Server Admin Engine tabs

The ARCserve Server Admin has four informational tabs:

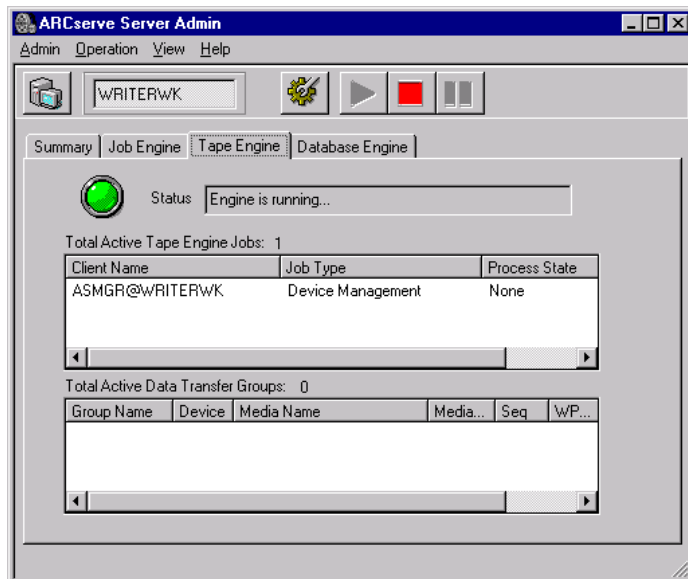
- Summary – shows you the current state of all the ARCserve engines and the number of days, hours, and minutes each engine has been running.



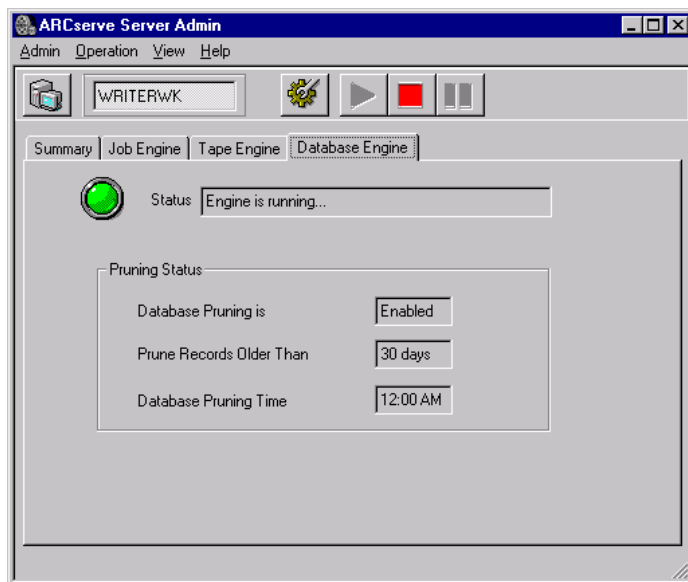
- Job Engine – displays information about the jobs submitted, such as the total number of jobs and the number of active, ready, hold, and done jobs. It also shows the queues, which are being scanned and the scanning interval.



- Tape Engine – displays information about jobs using the tape engine, such as the type of job it is and who submitted it. It also displays information on media groups.



- Database Engine – displays the database pruning information selected.



Engine Status

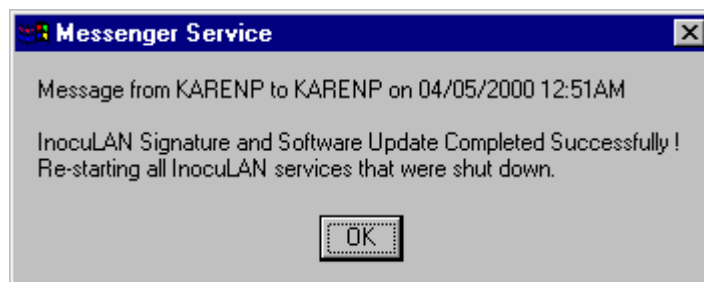
A stopped engine is an engine that is completely offline. This may be caused by errors, manual shutdown, or a new installation. Whatever the reason, it means that the services of that engine are not available.

ARCserve's engines are designed to run independently of each other. If you stop the Tape Engine, the Database Engine and the Job Engine are not affected. They continue to run, performing their services as configured. The Database Engine continues to log pertinent ARCserve information in the database, and the Job Engine continues to scan the job queue and launch jobs as required. If a job requires a storage device, the job is launched by the Job Engine as usual, but the job will fail because the Tape Engine is not able to communicate with the storage device. The Database Engine then logs this information.

Note: Although ARCserve still functions if one or two engines are not running, ARCserve needs all three engines running simultaneously to achieve complete functionality.

InoculateIT

Since InoculateIT, Computer Associates premier anti-virus protection solution is bundled with ARCserve, the job engine will need to be stopped and restarted under certain conditions to fully update the anti-virus protection. InoculateIT can be configured to automatically download updated virus signature files and program modules. These updates are then distributed to the participating applications. Once this has been accomplished, InoculateIT will broadcast a message stating that the update has been completed. At this point, the ARCserve job engine must be stopped and restarted.

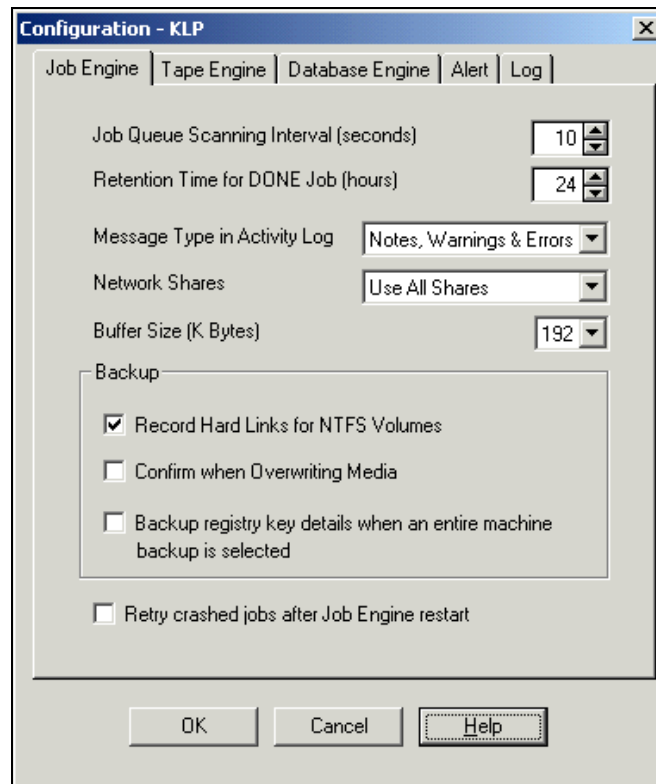


Configuring the ARCserve Engines

The ARCserve Server Admin allows you to configure each engine to suit your needs. To configure the ARCserve engines, select Configuration from the ARCserve Server Admin menu. Your configuration options include the following:

Job Engine Configuration

The ARCserve Job Engine controls the execution time of jobs in the job queue. It scans the job queue regularly, launching jobs as their execution dates and times are reached.



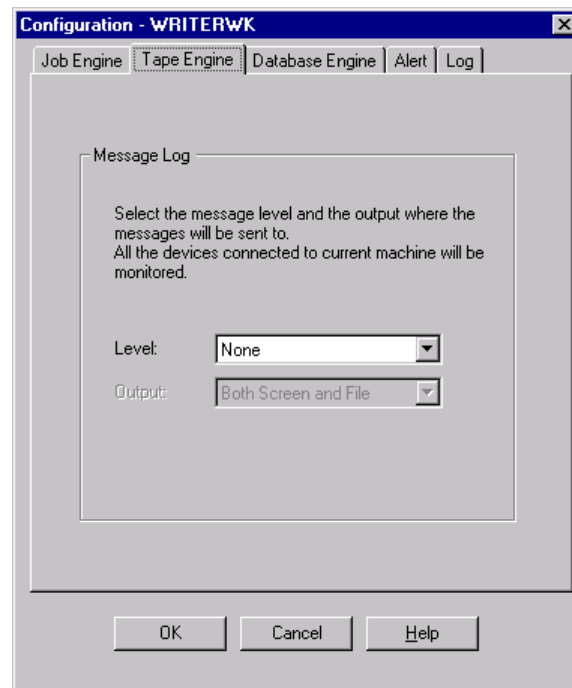
You can configure the following options:

- Job Queue Scanning Interval (seconds) – The Job Engine constantly scans the job queue for jobs that should execute. By default, the job queue is scanned every 10 seconds. To change the time interval, specify a number from 1 - 999.
- Retention Time for DONE Job (hours) – Jobs with a final status of “DONE” remain in the job queue for the time specified in this field. By default, ARCserve keeps DONE jobs for 24 hours before they are deleted from the queue. To change the time, specify a number between 0 and 999.

- **Message Type in Activity Log**— The Activity Log contains information about all of ARCserve’s activities. By default, notes, warnings, and errors that occur when running ARCserve appear in its Activity Log. To change the types of messages, specify one of the following values:
 - None: No messages appear.
 - Errors: Only errors that occur while running ARCserve appear.
 - Warnings & Errors: Warnings and errors that occur while running ARCserve appear.
 - Notes, Warnings & Errors: (default) Includes all notes, warnings, and errors that occur while running ARCserve.
 - Debug: Only debugging information appears. This should only be used for troubleshooting purposes.
- **Network Shares**— By default, ARCserve will display Use All Shares in the Browser. This means that both Default Shares and User Shares are available for selection as either your source or destination for a job. To change the type of shares that are displayed in the Browser, specify one of the following:
 - Use Default Shares Only: Only administrative shares are available.
 - Use Users Shares Only: Only shares that have been specifically set by users will be seen.
- **Buffer Size (K Bytes)**— Define the buffer size used by ARCserve.
- **Backup**— Allows you to customize additional options in your backup jobs:
 - Record Hard Links for NTFS Volumes: (default is ON) If you back up hard links files, this information is included and preserved by default.
 - Confirm when Overwriting Tapes: (default is OFF) Any time a tape will be overwritten, ARCserve can prompt you to confirm that you really want this to happen. By default, this option is not set. If you set this option on, ARCserve will pop up a confirmation dialog box and wait five minutes for you to respond. If you still have not confirmed after five minutes, the job will be canceled.
 - Registry key details when an entire machine is selected: (default is OFF) You can turn on the option to back up the Registry key details for target machines by clicking on the “Backup registry key details when an entire machine is selected” check box.
- **Retry Crashed Jobs after Job Engine Restart**— This option is a checkpoint mechanism. ARCserve will attempt to restart a crashed job if this box is checked. It should only be turned off if a cluster environment is configured to allow for fail-over.

Tape Engine
Configuration

The ARCserve Tape Engine identifies all the storage devices that are connected to your system and activates them when a job is started. Changing the configuration of the Tape Engine is not recommended. Keep the default configuration, except for troubleshooting purposes.



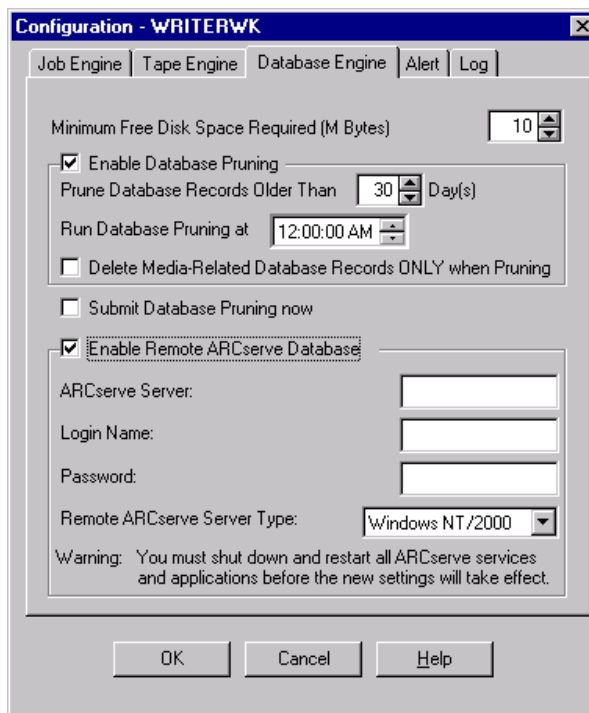
Your configuration options include:

- **Level** – If you keep the default ("None"), you do not need to specify anything else. The available values are:
 - None: (default) No information is logged. If you keep the default, the Tape Engine Log will not appear in the Job Status Manager.
 - Detail: Every message sent by the storage devices is logged. ARCserve will monitor the Tape Engine and log all information in a Tape Engine Log. The Tape Engine Log will appear in the Job Status Manager.
- **Output** – If you specified "Detail" above, you can specify where you want the messages sent. Specify one of the following:
 - Both Screen and File - The messages are recorded in the Tape Engine Log as well as to a DOS box (the Tape Engine Message window).
 - Screen Only - The messages are sent to the Tape Engine Message window only.
 - File Only - The messages are recorded in the Tape Engine Log only. You can view the Tape Engine Log in the Job Status Manager.

Database Engine Configuration

The ARCserve Database Engine determines the logging options for the database, which stores statistical information for all jobs processed. Some of the information saved is:

- What files and directories have been backed up, copied, and restored
- Jobs that ARCserve has processed
- Storage devices and media used for ARCserve operations



Your configuration options include:

- **Minimum Free Disk Space Required (M Bytes)** – Specify the minimum amount of disk space (in megabytes) you want to reserve for the ARCserve database. By default, 5 MB of disk space is specified. To change the amount of disk space required, specify a number between one and 10.
- **Enable Database Pruning** – Information regarding the files and directories that were backed up or copied in a session is deleted when database pruning is on. By default, this option is selected, in order to free up space in the database file. However, it can be useful to set this option off in order to maintain the detailed information for restoring purposes. However, not pruning your database can mean that it fills up quickly.

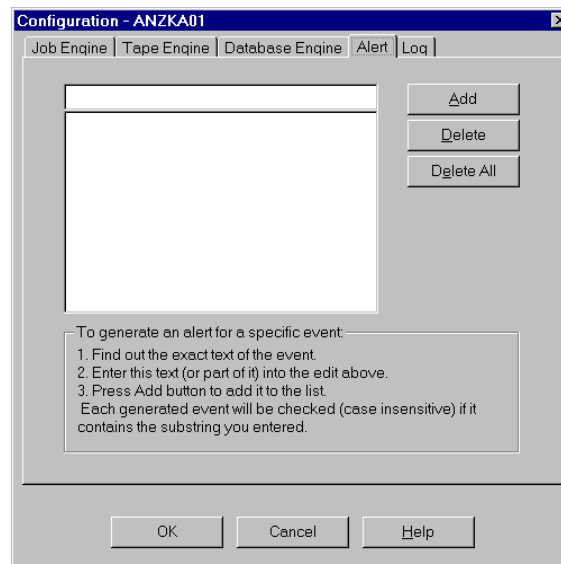
Prune Database Records Older Than _ Day(s) – This field is active only if the Enable Database Pruning option is on. Specify how long the record should be in the database before ARCserve deletes it. By default, the record will be deleted (if enabled) after it has been in the database for 30 days. To change the length of time, specify a number between one and 365.

Run Database Pruning at _ – This field is active only if the Enable Database Pruning option is on. Specify when you want the pruning operation to run. By default, pruning (if enabled) will occur at 12:00 AM.

- Delete Media-Related database Records ONLY when Pruning – When you reformat or erase a media, ARCserve will delete the records in the database that pertain to the tape. Performing this extra step, however, can be a time-consuming process. Select this option to postpone deleting these records until pruning is performed.
- Submit Database Pruning now – Submits a database pruning job to the ARCserve job queue with the specified parameters immediately.
- Enable Remote ARCserve Database – This enables the ARCserve Server to use a remote ARCserve database as its local database, regardless of the remote database type. This database can be used to submit a restore job to the local ARCserve machine. Configuration of the remote database is performed here as well. Enter the ARCserve Server Name, Login Name, Password and Server Operating System type of the remote ARCserve Database you want to enable.

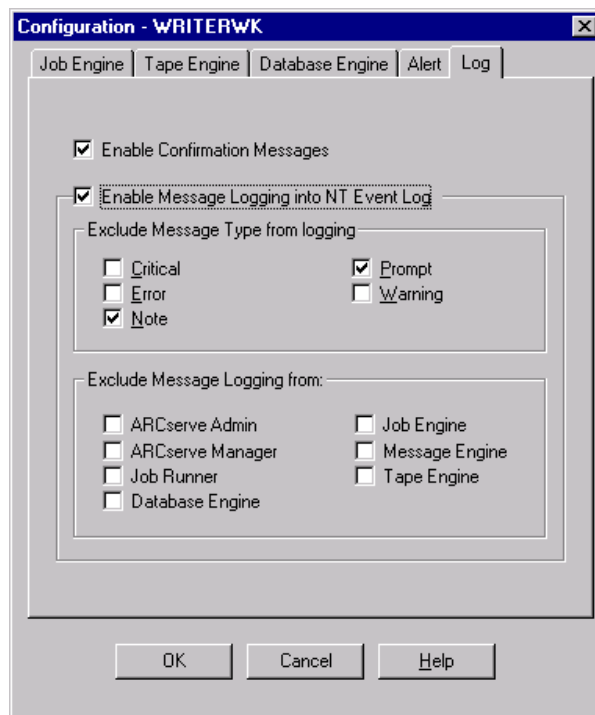
Alert Configuration

Configures the Alert notification for non-job related events, such as the tape engine starting or stopping, successful completion of operations. You can have the Alert program send an alert message if the ARCserve Server comes across a specified message from within the ARCserve logs. For more information on the Alert feature, refer to the chapter, “Using the Alert Manager”.



Log Configuration

This allows you to specify which messages may be written to the event log. By default, all messages are written.



If you enable Message Logging into NT Event Log, the following fields are enabled:

- Exclude Message Type from logging – You can select which type of message should be excluded from the event log: Critical, Error, Note, Prompt and/or Warning.
- Exclude Message Logging From – You can choose to exclude all messages from a particular ARCserve module: ARCserve Admin, ARCserve Manager, Job Runner, Database Engine, Job Engine, Message Engine and/or Tape Engine.

Additional Server Admin Functions

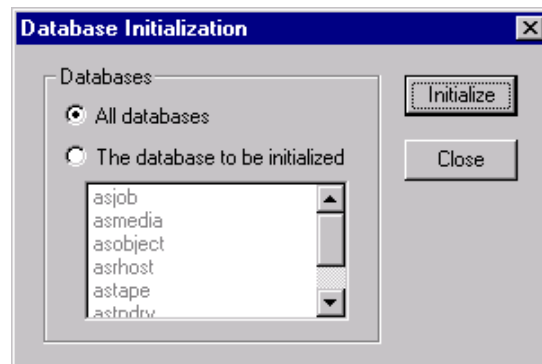
Changing the
ARCserve System
Account

The ARCserve Server requires a valid user account on the host Windows NT/2000 machine (initially entered during installation). You can change the system account at any time in the ARCserve Server Admin program.



Initialize Database

You can initialize a particular ARCserve database by selecting Initialize Database from the Operation menu. This option is only applicable after selecting the database Engine tab.



Using the Alert Manager

Alert is a notification system that sends messages to people in your organization using various methods of communication. Alerts can be sent to the system administrator, a hardware technician, or anyone else, in or out of the office. An individual or groups of persons in different segments of the network can also be notified.

In order for ARCserve to generate alerts you must provide Alert with the necessary information to communicate. For example, if you will be using the pager system, you will have to provide the pager number to dial, and you will have to supply information about your modem. All of this information must be configured in the Alert program on your server.

Alert does not generate its own messages. Alert routes any messages, errors, or warnings it receives from different sources including ARCserve and distributes these “alerts” to specific destinations. For example, ARCserve generates warning messages whenever a virus is detected. These warning messages are passed to Alert, which sends the notification. Alerts can be sent via:

- Broadcasts – Alert broadcasts can be sent to specific computers.
- CA-Unicenter TNG Option – Sends a message to the TNG console and/or WorldView repository when an alert is generated.
- Local and remote NT Event Log Notification.
- Electronic Mail – Microsoft Exchange and Lotus Notes.

Tip: You can use either the Alert Setup program or the Exchange Mail Setup option in the Alert Manager File menu to configure MS Exchange Mail for use with Alert.

- Pager – Alphanumeric only.
- Simple Network Management Protocol (SNMP) managers – Such as NetWare Management System (NMS), HP OpenView, and CA Unicenter TNG.
- Trouble Tickets – An alert can be printed through any print queue on your network.

Components

- Alert Service— This is the service that is responsible for the reception, processing, and distribution of Alert messages.
- ALBUILD.DLL— This is the .DLL that acts as the channel between Alert and other applications. This should be located in the Alert home directory.
- *.ALT— This is the application profile file. This file is provided by an application (such as ARCserve, which creates an ARCSERVE.ALT file). This .ALT file must be present in the Windows 2000 directory in order for Alert to handle messages generated by an application.
- Alert Manager— This is where you can configure how Alert sends its messages.

Alert Features

ARCserve provides remote management and configuration of the Alert Service. Alerts from clients can be sent using IP in addition to the standard IPX protocol. Messages can be sent containing full paths of any virus-ridden files.

The Alert Manager

You can run Alert from the ARCserve program group by selecting Alert Manager. You can use the Alert Manager to select a remote machine to manage Alert messages. Before you start Alert, you must establish a Service Account connection and select a remote machine.

Note: If you previously had the Alert Manager installed with another Computer Associates product, it will not re-install the Alert Manager to the ARCserve directory, it will remain in the directory where it was previously installed.

Configuring Alert

Alert allows you to configure default settings used by all applications that use the Alert Service. You can also enter configuration information specifically for an individual application, which will override the default Alert configuration. Each application that uses Alert is displayed as a leaf on the left-hand function tree.

Creating and Editing Port Configurations

The Ports object, located under the Configuration object, contains communication port profiles. The following Port configurations are for any function that uses serial port access:

- Port—The name of the communications port you want the pager message to be broadcast from.
- Baud Rate—The baud rate being used by your modem.
- Parity—The parity setting, none, odd, or even, of your modem.
- Data Bits —The number of data bits, 7 or 8, that your modem uses.
- Stop Bits —The number of stop bits, 1 or 2, that your modem uses.

Alert's Broadcast Option

Alert broadcasts can be sent to specific network users or groups. To learn how to add broadcast recipients, refer to the online *ARCserve Procedures*.

Alert's Pager Option

The pager option is used to send an alphanumeric pager message. When you highlight the Pager option, the current list of recipients appears. To learn how to add pager recipients, refer to the online *ARCserve Procedures*.

Tip: Before you can add pager recipients, you need to configure your communications ports.

Note: When sending an alphanumeric page, consult your pager manual for proper modem settings.

Interpreting the Pager Message

There are several messages similar to the ones below that can be sent to an alphanumeric pager. Words that appear in brackets will be filled in with an actual user name, workstation address, path and file name, virus name, or server name.

- Boot Virus Detected — username at workstation address
- Manager Detected a Virus [virusname] in [path] — username at workstation address
- Infected File [servername/path] Detected
- Infected File [path] Accessed by username at workstation address

SNMP Option

You can use the SNMP option to send an SNMP 'trap' (message) to an SNMP manager. Examples of SNMP managers include NetWare Management System (NMS), HP OpenView, IBM NetView and CA Unicenter TNG.

The online *ARCserve Procedures* explains the SNMP fields on the SNMP Configuration screen, and how to use them.

Trouble Ticket Option

Trouble tickets are used to alert users through a printed document.

E-mail Option

The E-mail option is used to send e-mail messages to specific users.

Important! *Lotus Notes or Microsoft Exchange Client must be installed on your computer in order to be able to send messages or to enter configuration data on this screen. Consult your Windows NT/2000 manual for instructions about how to set up your e-mail account.*

Unicenter TNG Option

The Unicenter TNG (The Next Generation) Option makes it possible to send a message to the Unicenter TNG console and/or World View repository when an alert is generated.

Note: The Alert application must be running on both the Event Management machine and the WorldView machine.

Refer to the online *ARCserve Procedures* for explicit directions on how to send a message to the Unicenter TNG Console and/or the WorldView repository.

Sample TNG Alert Scenarios

If you want to send informational alerts to the Unicenter TNG Console using blue text, for example, configure a recipient as follows:

Event Priority	Description
Informational	Application Event Priority
Blue	Color
✓	Send to Console
✓	Send to World View

If you want to send error alerts to the Unicenter TNG Console using red text, and have the object status in the World View repository updated, configure another recipient as follows:

Event Priority	Description
Critical	Application Event Priority
Red	Color
✓	Send to Console
✓	Send to World View

Event Log Option

You can configure the Event Log so that Alert will put an event for a selected server in that machine's Event Log.

Event Priority

All applications calling Alert specify one of the following Event Priorities:

- Critical
- Warning
- Informational

Alert's Activity and Event Logs

While the current status of Alert is shown when the Alert Summary is highlighted in the Activity group, a historical listing is stored in the Activity Log. Similarly, every message that is generated by Alert is stored in the Event Log. You can view, print, or clear these logs. See the online *ARCserve Procedures* for directions.

Testing the Recipients

You can use the Test toolbar button to test any of the Alert messaging functions without an actual “alarm” condition. Refer to the online *ARCserve Procedures* for directions on how to do this.

Note: You should test any features after the configuration has been completed. Be sure to inform any Alert recipients that a test is taking place.

ARCserve Cluster Support

Microsoft Cluster Service (MSCS) allows multiple Windows NT/2000 operating system-based servers to be connected together, making them appear to network clients as a single, highly available system. This appendix provides a high-level overview of ARCserve cluster support, which is capable of fail-over as well as backing up, protecting and recovering the cluster node.

The initial goal of the ARCserve cluster support was to protect cluster nodes from disasters by enabling ARCserve to backup/restore the cluster specific resources, such as shared disk, quorum resource, disk signature, and the cluster registry hive. The main goal of the ARCserve version 7.0 cluster support is to be a more cluster-ware application and to be a fault-tolerant application that can handle fail-over in a cluster environment and in a non-cluster environment as well.

ARCserve improves availability, scalability and manageability through the clustering technology.

- **Availability through fail-over:** When an ARCserve server in the cluster fails, the ARCserve services can recover and disperse the ARCserve job from the failed server to another server in the cluster. As a result, the failure of an ARCserve server in the cluster will not affect high service availability for the users.
- **Scalability through Active/Active configuration:** When the overall load exceeds the capabilities of the systems in the cluster, by adding additional systems to the cluster and by balancing the overall load among the ARCserve in the cluster, the ARCserve can allow customers to incrementally add smaller, standard systems, as needed, to meet overall backup and restore processing power requirements.
- **Manageability through Cluster Resource DLL:** The ARCserve in the cluster extends cluster service capabilities to monitor and detect failures of its own components. The network name and/or IP address of a virtual server enables ARCserve to appear as a single system to end-users and takes advantage of the existing central management capabilities of cluster management so that the nondedicated support staff can configure and maintain the ARCserve in the cluster.

Cluster-aware Application	The ARCserve has its own Cluster resource DLL to extend Cluster Service capabilities to monitor and detect failures of ARCserve. The ARCserve services can be brought online or taken offline gracefully.
ARCserve Services Fail-over Capability	Through its own Cluster resource DLLs, Cluster Service can start restart ARCserve automatically from the other cluster node in case of failure. Finally, the most important feature is to support Active/Passive fail-over capability. With Active/Passive capability, a copy of the ARCserve resource type can be run on each node in the cluster, and by telling the ARCserve resource on each node to release and acquire the ARCserve operation, as appropriate, the ARCserve services can be moved from node 1 to node 2.
ARCserve Job Fail-over Capability	When the ARCserve job fails due to the failure of a node in the cluster, the cluster service will recover ARCserve services from the failed node to another node within the cluster. Once the ARCserve services are recovered from another node, the failed ARCserve job will rerun from the recovered node.

Requirements for Installing ARCserve on a Cluster

The installation requirements can be grouped into the following two categories:

Software Requirements

- The ARCserve system accounts must be the same among all ARCserve servers installed in cluster nodes.
- ARCserve must be installed in each cluster node with the same software and hardware configuration.
- In case of Active/Passive configuration, ARCserve must be installed into the same lettered and shared drive/volume in order to handle job fail-over.

SCSI Environment

- Base ARCserve and the Tape Library Option (if libraries are used).
- Use the default group names, or make sure that the same group names are used on ARCserve on both the nodes.
- Duplicate "Barcode" key under the tape engine, using cluster service.

Fibre Environment

Add the following registry switch under the Tapeengine Key:

Add Key: CONFIG

Add Value : EnableSharedDevices: DWORD : 1

Hardware Requirements

SCSI Environment

- Both cluster nodes should have identical hardware configurations (SCSI adapters, hard disks etc.)
- It is recommended that separate SCSI controllers are used for connecting tape drives/tape libraries.

Fibre Environment

- Both cluster nodes should have identical hardware configurations (Fibre channel controllers, hard disks etc.)

Note: It is recommended that all hardware used be similar, if not identical, for both nodes. This will make configuration easier and will eliminate any potential compatibility problems.

Installing ARCserve on a Cluster

To enable ARCserve services and jobs fail-over, the ARCserve must be installed on the same drive letter cluster shared disk/volume on each cluster node.

1. Install ARCserve on node 1.
2. Install Tape Library Option on node 1.
3. Configure the Shared Device on node 1.
4. Install ARCserve on node 2.
5. Install Tape Library Option on node 2.
6. Configure the Shared Device on node 2.
7. Create a new Registry Key with DWord Value of 1 on node 1.
8. Create a new Registry Key with DWord Value of 1 on node 2.
9. Create the ARCserve resource and group through the command prompt, and register the ARCserve system registry key for replication.
10. Finally create a file share resource to ensure that ARCserve is shared after fail-over.

Registering ARCserve with the Cluster

During the ARCserve installation, the setup program copies `ascluster.dll` and `asclusterex.dll` to “%WINNT_ROOT%\system32” directory.

Note: The following commands must be entered exactly as shown to properly register the product with the cluster.

To add the ARCserve resource type to the cluster, use `cluster.exe`:

```
cluster resourcetype "ARCserve" /create /dllname:ascluster.dll
```

The `cluster.exe` utility can be used to create resources of type "ARCserve". To create a resource using `cluster.exe`, use a command like the following:

```
cluster GROUP "My ARCserve Group" /create
cluster resource "My ARCserve" /create /group:"My ARCserve Group"
    /type:"ARCserve"
```

When the ARCserve is installed into a cluster disk, we recommend you combine the “ARCserve” and the “Physical Disk” resources into the same group so that, in case of fail-over, all of the resources in the group fail-over together as one unit to another node in the cluster.

To guarantee the consistency of the data stored in the system registry, the ARCserve registry key can be registered for replication. Then the Cluster Service updates the system registry on the current node in case of fail-over from the other node.

From the cluster administrator in the active cluster node with the ARCserve Group online:

1. Add a new resource to “My ARCserve Group” with the following entries:
Resource Name: “My ARCserve Registry”
Resource Type: “Generic Service”
Group: “My ARCserve Group”
2. Specify the generic service name as “astapeengine”
3. Add a Registry Key “SOFTWARE\ComputerAssociates\ARCserveIT\Base” to the replication key.

To guarantee that the ARCserve directory is shared to ensure Fail-Over reliability the ARCserve directory should be added as a resource to the ARCserve Group. From the Cluster Administrator add the following resource to " My ARCserve Group" .

1. Add a new resource to "My ARCserve Group" with the following entries:

Resource Name: "My ARCserve Shared"


Resource Type: "File Share"

Group: "My ARCserve Group"

Share Name: "ARCserve\$"

Path: "S:\Program Files\ComputerAssociates\ARCserve"

Permissions: "Administrator"



To uninstall or delete ARCserve resource from the cluster, make sure that ARCserve resource is offline and do the following:

```
cluster resourcetype "ARCserve" /delete
cluster resource "My ARCserve" /delete
cluster GROUP "My ARCserve Group" /delete
```

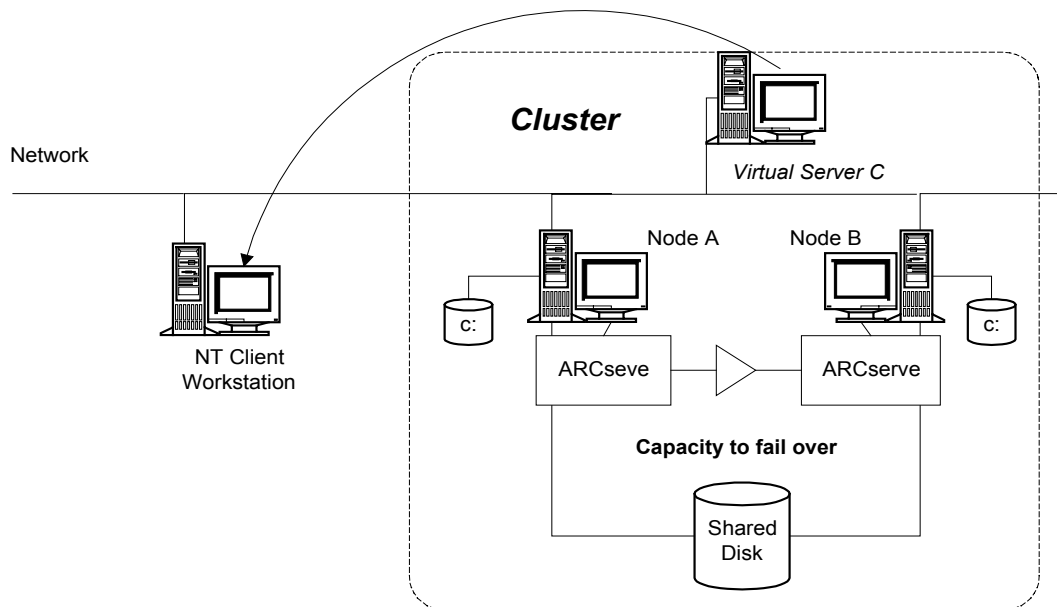
ARCserve Services Fail-Over (Active/Active & Active/Passive)

Once ARCserve is registered with the cluster, the cluster service provides ARCserve services for automatic fail-over. This fail-over solution improves ARCserve availability by allowing two or more servers to share the same virtual cluster node name and the same ARCserve installation on a shared hard disk within a cluster. The ARCserve services fail-over is supported in both Active/Active and Active/Passive configuration.

ARCserve Job Fail-Over (Active/Passive)

When an ARCserve server in the cluster fails, the cluster-aware ARCserve will recover and disperse the ARCserve job from the failed server to another server within the cluster. Once the ARCserve services are resumed in another cluster node, any ARCserve job failed from the previous server will be rerun in a new cluster node. Only in Active/Passive configuration, ARCserve job fail-over is supported because only in this configuration, the ARCserve job queue and database can be shared among the nodes in a cluster.

All the ARCserve system accounts across the cluster nodes should be the same, and the job submission to the ARCserve installed in the cluster should use the virtual cluster node name as its ARCserve host name instead of its physical node name.



When ARCserve is failed over to the other cluster node, ARCserve will recover and pick up the failed job from the ARCserve queue and resume the job from the volume where it failed.

ARCserve Checkpoints

When a failed ARCserve job is rerun in a new cluster node, through the Checkpoints mechanism, the ARCserve job can be restarted at the volume/disk where it failed.

Index

A

Activity Log, 11-6

Alert

 Basic Components, 11-2

 Broadcast, 11-3

 Configuring, 11-2

 Event Log, 11-6

 Pager, 11-3

Alert notification system, 1-3

ARCserve Engines

 Configuring, 10-5

 Viewing and Changing the Engine's Status, 10-4

ARCserve Server Admin, 1-5

ARCserve Wizards, 1-8

B

First Media options

 First media options, 2-4

Session password

 Session password, 2-5

Backup Database Agents, 2-1, 3-2

Backup Job, 4-7

Backup Manager, 1-7

Backup Wizard, 2-1

C

Calendar View

 Customizing GFS (Grandfather-Father-Son)Rotations, 7-11

Client agent, 1-2

Copy Operation Options, 1-13

D

Database Manager, 1-7

Device configuration, 1-3

Device Group, 6-7

Device groups, 6-7

 Configuring, 6-7

Device Management, 1-7

Device Manager

 Functions, 6-2

 General information, 6-1

 Getting information, 6-7

E

Erase Options

 Long Erase, 6-4

 Quick Erase, 6-4

 Quick Erase Plus, 6-4

Event Log, 11-5, 11-6

F

File Pattern filter, 5-2

Filters

- Directory Filter, 5-2
- Exclude Over Include, 5-1
- File Attributes Filter, 5-3
- File Modified/Created/Accessed Filters, 5-3
- File Pattern Filter, 5-2

Filters, using, 5-1

G

GFS rotation

- Media pools, 7-8

J

Job Log, 4-5

Job Queue Menu, 4-3

Job script, 5-6

Job Status Manager, 1-7

M

Media and Device Management

- Media spanning, 1-2

Media options

- Additional media options, 2-5
- Backup
 - Media options, 3-7

Media Overwrite/ Append Rules, 2-3

Media Pool, 7-1

- Assigning Media, 7-4
- GFS Media Naming, 7-7
- GFS Rotation Schemes, 7-6
- Grandfather-Father-Son (GFS) Rotation, 7-5
- Retention period, 7-2
- Save Set, 7-2
- Scratch Set, 7-2
- Using a Simple Rotation Scheme, 7-6

Media Pool Management, 1-8

Media pooling

- GFS media pools, 7-8
- Location maintenance, 7-4

Media Rotation Scheme, 1-2

Media spanning, 6-7

Media, Merging, 1-11

P

Pager

- Testing, 11-6

Parallel streaming, 6-7

Pre/post Commands

- Specifying for Backup Job, 2-9

R

RAID (Redundant Array of Inexpensive Devices), 1-2

Removable Drive Support, 6-8

Restore jobs

- Source views
 - About, 3-3

Restore Wizard, 3-1

Retention time, 7-9

S

Save set, 7-9

Scheduling a Job, 5-5

Scratch set, 7-9

Serial number information, 7-5

- Automatic naming, 7-5
- Bar Code, 7-5

Shares

- Administrative;Administrative shares;, 2-1
- User specified;, 2-1

T

Tape Engine, 6-7

Trouble Ticket, 11-4

U

Use Rotation Scheme

 Selecting GFS (Grandfather-Father-Son) Rotation
 Schemes, 7-11

Using Unicenter TNG Option

 With Alert, 11-4

V

Verifying Data, 2-6