

ARCserve for Windows NT v2.01 User Guide

©Copyright 1995, Cheyenne Software, Inc. All Rights Reserved.

U.S. GOVERNMENT RESTRICTED RIGHTS

The software and accompanying materials are provided with RESTRICTED RIGHTS. Use, duplication or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 or the Commercial Computer Software Restricted Rights clause at FAR 52.227-19 subdivision (c)(1) and (2), as applicable. Contractor/manufacturer is Cheyenne Software, Inc., 3 Expressway Plaza, Roslyn Heights, New York 11577.

Cheyenne Software Inc. provides this publication “as is” without warranty of any kind, either expressed or implied, including but not limited to the implied warranties of merchantability or fitness for a particular purpose. The entire risk as to the use of this information is assumed by the user.

In no event will Cheyenne be liable for any damages, direct, indirect, incidental, special or consequential, resulting from any defect in the information, even if it has been advised of the possibility of such damages.

Further, Cheyenne reserves the right to revise this publication and to make changes to it from time to time without obligation to notify any person or organization of such revision or change.

Trademarks

Cheyenne and ARCserve are registered trademarks of Cheyenne Software, Inc.

Other brand or product names used in this manual, but not listed here, are trademarks or registered trademarks of their respective holders.

Preface

This manual is written for people who will use ARCserve. It contains information about making a backup, copying files, restoring files, managing files in the job queue, managing your tapes and tape drives, viewing information in the ARCserve database, viewing and printing reports, and using ARCserve's utilities.

Credits

Written by William Curtin and Chris Merlo

Product Support

If you have any questions regarding the use of this product, please contact us at one of the following locations:

USA, Canada, Asia, Latin America:		
3 Expressway Plaza Roslyn Heights New York, 11577	Main Number	516-484-5110
	Technical Support	516-484-3493
	FAX	
	Technical Support	800-CHEY-TEC Mon-Fri 8 am-6:30 pm EST Mon-Fri 6:30 pm-10 pm EST (Callback only) Sat/Sun 10 am-4 pm EST (Callback only)
	BBS	516-484-3445
CompuServe	GO CHEYENNE	
<hr/>		
France (European Headquarters):		
Cheyenne Software S.A.R.L. Bel Air Building, 58 Rue Pottier 78150 Le Chesnay, France	Tech Support	+33-1-39 23 18 70
	Hotline	
	Tech Support (FAX Hot Line)	+33-1-39 23 18 69
	BBS	+33-1-39 23 18 60
<hr/>		
Japan:		
Cheyenne Software K.K. Chiyoda Bliss Building, 5F 2-24, Kanda Sakuma-cho Chiyoda-ku, Tokyo 101 Japan	Voice	(03)3865-0191
	FAX	(03)5820-4781

1

Chapter

ABOUT ARCSERVE FOR WINDOWS NT

ARCserve for Windows NT is your solution to data security and management for your Windows NT network. Providing the broadest and highest-level features available, ARCserve protects you and your network from disaster.

In this chapter you will learn:

Page

- 1-2 ➤ About the Enterprise and Single Server versions
- 1-4 ➤ About the key features provided by ARCserve
- 1-6 ➤ How ARCserve works
- 1-14 ➤ Style conventions used in this manual

What is ARCserve for Windows NT?

ARCserve for Windows NT is a data management and backup program designed to run on your Windows NT machine.

Two versions of
ARCserve are
available

ARCserve for Windows NT is available in two versions, the ARCserve Enterprise version and the Single Server version. With the Enterprise edition of ARCserve for Windows NT, you can perform immediate, scheduled, or automated backups of your local and remote Windows NT machines attached to your network.

The Single Server version is able to back up only the host server and any additional share points that are mapped to that server. However, both ARCserve for Windows NT versions can work together. For example:

The Boston headquarters of XYZ Transportation has installed the Enterprise version of ARCserve for Windows NT. The company maintains branch offices in Chicago and Houston, with each local site running Single Server ARCserve for Windows NT and attaching by network to Boston.

Harry, the Boston Enterprise administrator, can monitor the branch office's jobs on the Global Job Queue. If the Chicago manager is away from the office, Harry can manage all of the Chicago activity.

The following table compares critical features of ARCserve for Windows NT Enterprise edition and Single Server edition:



Feature	Single Server version	Enterprise version
Remote management of ARCserve Servers	Allows you to back up the host server and any share points that are mapped to the host	Allows you to manage multiple servers from a single machine. Gives you the ability to back up, copy, and restore any Windows NT machine and Novell NetWare server in your network
Enterprise client agent support	Not available in the Single Server version.	Lets you use Cheyenne client agents (such as Client Agent for NetWare), which employ the latest PUSH technology to transfer filtered data to the host server for processing. This enables the host server to be more efficient and to process faster




Buttons, menu options, etc., that appear “grayed out” (non-selectable) in ARCserve for Windows NT Single Server version are features that are not supported.

Software preview

ARCserve for Windows NT high performance features

- “Parallel streaming” - with multiple tape devices, you can have multiple jobs running at the same time.
- Tape spanning - large jobs requiring more than one tape automatically continue (span) to the tape(s) that you provide.
- Tape rotation - get the most accurate and complete backups of your data by scheduling repeating jobs that use different method options to back up files.
- Single Server - reduces network traffic because the data can stay locally.
- Remote Management of ARCserve Servers - the Enterprise version allows you to manage multiple servers from a single machine. Gives you the ability to back up, copy, and restore any Windows NT machine and Novell NetWare servers in your network.
- InocuLAN virus scanning - lets you scan files during a backup or copy operation and notify users if a virus is detected.
- Event-driven notification system - the Cheyenne Alert notification system informs you of key events in your operations.
- Enterprise client agent support - lets you use Cheyenne client agents (such as Client Agent for NetWare), which employ the latest PUSH technology.
- Windows NT 3.51 - supports compressed files, directories, and drives on Windows NT File System (NTFS) formatted drives. Also supports Service Security Impersonation.



- Windows-based Graphical User Interface - uses all the familiar Windows features including drop-down menus, descriptive buttons and icons, and dialog boxes.
- Quick Status information - allows you to find information quickly.
- User-defined scripts - allow you to configure jobs once and re-use them when needed.
- Various database views - give you quick access to information about jobs performed by ARCserve and help you quickly locate and restore your data.
- Reports - give you a complete history of jobs performed and specific activity on the machines, directories, and tapes that ARCserve has used.

Flexibility

- Intelligent restore - provides you with four methods by which to select the data you need to restore.
- Support for a wide variety of SCSI 4mm, 8mm, DLT, and QIC tape drives.
- Auto Changer (optional) - allows you to use an autochanger with your backup rotation.
- Ability to back up, copy, and restore any Windows NT machine and Novell NetWare server in your network.

How does ARCserve work?

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

- The ARCserve Manager
- The ARCserve Server

The ARCserve Manager

The ARCserve Manager enables you to submit job requests to an ARCserve Server and monitor the status of your requests. ARCserve Manager lets you remotely manage multiple ARCserve Servers from a single machine in the following manner: you can

- Open a Manager window, such as the Backup window or the Restore window, for each job request.
- Easily select an ARCserve Server using the Select Server button on the Manager window.
- Manage all backup operations for the selected ARCserve server.
- Open a new window for each additional server that you wish to manage.

- If you are an Enterprise user, you can switch between job requests on different servers by changing your Window focus.

For example, you can open a Backup window, select server A and schedule a backup for this server. Then open a Restore window, select server B and schedule a restore for server B. Now switch to your Backup window, and check on the status of your backup job and server A. Now open a second backup window select server C and so on.

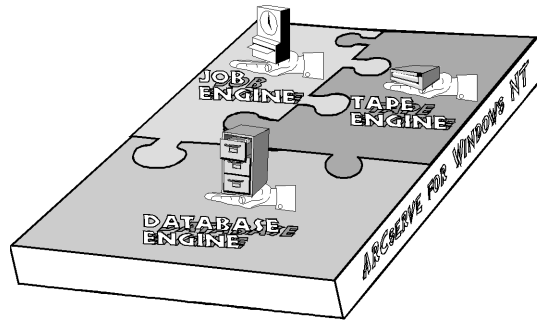
1

The ARCserve Server

Using the information you specified in an ARCserve Manager window, the ARCserve Server is the piece that performs the operations you set up. For example, suppose you used the Backup Manager to set up a backup job, selecting a source, destination, and date and time for it to run. The ARCserve Server takes this information and backs up the files to the destination at the time you specified. The ARCserve Server consists of three engines. These engines are the:

- Job Engine
- Tape Engine
- Database Engine

Though the three ARCserve engines are separate pieces, together they make the ARCserve Server.



For peak performance, all three engines must be running

Each engine works with the others to back up, copy, restore, and perform all the ARCserve jobs you submitted through the ARCserve Manager. To achieve ARCserve's fullest potential, all three engines should be running concurrently.



Any time we refer to the "ARCserve Server" in this manual, we are referring to the ARCserve Job Engine, ARCserve Tape Engine, and ARCserve Database Engine as a unit.

The Job Engine

The Job Engine is responsible for processing your jobs at their designated date and time. It constantly scans the job queue. When it finds a job that is ready to run, the job engine sends it to the appropriate handler.

The Tape Engine

The Tape Engine is responsible for communicating with and controlling your tape devices. When a tape device is needed for an ARCserve job, the Tape Engine lets your hardware know it is needed.

The Database Engine

This engine is responsible for maintaining the history of several things:

- Files, directories, drives, and machines that ARCserve has backed up or copied.

- Information about jobs that have been processed by ARCserve, such as the type, the final result, the start and end time, etc.
- Tapes used by ARCserve, such as the name, date it was first formatted, date it expires, sessions on it, etc.

1

By default, everything that happens within ARCserve is recorded by the Database Engine.

Software Preview

This section briefly describes ARCserve's features.

Quick Status

Provides a quick and convenient way of keeping track of ARCserve activity for each operation you perform. From this Manager, you can view the Activity Log for information. If you configured the Tape Engine (with the ARCserve Server Admin) to monitor your tape devices, you will also see the Tape Engine Log here.

Backup

The Backup Manager enables you to back up data to tape. ARCserve performs a complete backup of Windows NT machines. All Windows-NT specific information, such as registry information for each machine, is backed up intact.


The optional Client Agent for Netware is available for backing up NetWare file servers. Without the agent, NetWare-specific files (such as bindery and NDS information) are backed up as DOS files. With the agent, all NetWare-specific files are backed up intact, ensuring complete backup/restore of the NetWare file server. Client agent software is also available for other operating systems.

Information about each backup job (such as the path and name of each file, as well as the tape used) is logged in the ARCserve database.

Restore

The Restore Manager gives you four different approaches to restoring your data. They are:

- **Restore by Tree** - Select from a tree of machines that have been backed up.
- **Restore by Tape** - Select from a tree of tapes that have been used for backups.



- **Restore by Query** - Search ARCserve's database for files and directories that have been backed up.
- **Restore by Tape Media** - Select a session from a tape currently in the tape drive.

Copy

Use the Copy Manager to copy files and directories from one drive to another. This feature is especially useful for creating mirrored machines, drives, or directories.

Job Queue

Jobs that you have sent to ARCserve for processing can be managed from the Job Queue Manager. All pending, completed, and active jobs are listed in this Manager window. "Done" and pending jobs can be rescheduled, and active jobs can be monitored here. Subsequent log information is provided after the job is finished. In addition, you can submit new jobs, delete pending jobs, and stop active jobs from this Manager.

Device

ARCserve's Device Manager allows you to view information about your tape devices and tapes, change a tape drive's compression mode, configure and manage tape groups, and perform tape maintenance functions such as formatting, erasing, and retensioning.

Database

The Database Manager provides information about the ARCserve database. It offers a series of different viewing methods for the ARCserve database. Depending on the view you select, the Database Manager window displays information such as the jobs processed by ARCserve, the tapes used by ARCserve, and the devices you are using with ARCserve. The information contained here is constantly updated and used by ARCserve.

Reports

The Reports Manager generates and displays reports based on the ARCserve database. For example, the Tape Media report helps you keep track of tape usage so that you can determine when a tape should be retired.

Utilities

Several Utility Manager windows allow you to perform useful operations:

- **Recover Database** - overwrite your existing ARCserve database with database information from a tape.
- **Merge Tape** - add information from a tape to your existing ARCserve database.
- **Scan Tape** - report information about one or more tape backup sessions.
- **Compare** - compare files on a tape to their counterparts on a machine.
- **Count** - count files on a machine.
- **Purge** - delete files from a machine.

Finding information in this manual

Scanning the left column of each page

You'll notice that this manual uses a two-column layout. In the left column of each page, you will find three pieces of information:

- Margin headings
- Figure captions
- Note and warning icons

You can scan these items and get a summary of the information that is presented on each page.

The margin headings serve two purposes: they divide the page into sections (normal header function), and they act as a short summary for the section (paragraphs) that they represent.

The figure captions also serve two purposes: they identify a figure, and they provide additional information about concepts being discussed.

Notes and warnings each have an associated icon that appears in the left column. You can scan for these icons for information that is critical to using ARCserve.

This is the icon for a note. Its purpose is to remind you to do something or to emphasize some important information.



This is the warning icon. Its purpose is to draw your attention to critical information.

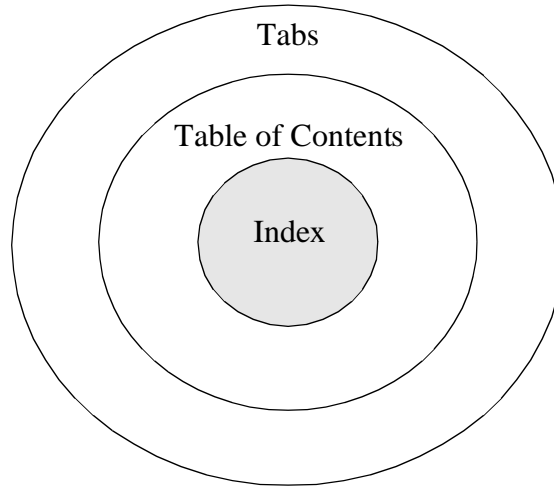


Using the tabs,
Table of
Contents, and
Index

In addition to scanning the left column on each page, you can also use the tabs, Table of Contents, and Index to search for specific information.

Use the tabs to find the chapter you want. Use the Table of Contents to find specific sections within a chapter. Use the Index to search for information by specific words or phrases, such as “Quick Erase”.

Think of using the tabs, Table of Contents, and Index as aiming at a target. Tabs get you close, the Table of Contents gets you closer, and the Index should get you a bull's-eye.





2

Chapter

GETTING STARTED

This chapter explains how to install ARCserve on your machine. It also guides you through a quick backup of your machine.

In this chapter, you will learn how to:

Page

- 2-4 ➤ Install ARCserve on your machine
- 2-12 ➤ Use tape devices that were installed through *Windows NT Setup*
- 2-14 ➤ Start ARCserve
- 2-15 ➤ Back up your machine - a tutorial

Checklist for getting started with ARCserve

The following checklist summarizes the steps you need to take to install and start using ARCserve. Each row in the table includes one general action that must be taken, followed by some helpful hints and suggestions, and a page or chapter reference for more detailed information.

Action	For more information refer to:
Make sure your machine has at least one SCSI board supported by Windows NT and by ARCserve -ARCserve supports up to 12 installed SCSI boards	For a list of certified SCSI boards, refer to Cheyenne's list of approved devices.
Attach a certified tape drive -You must make sure the tape drive is initialized before turning on your machine, otherwise Windows NT and ARCserve will not register its presence.	For a list of certified tape drives, refer to Cheyenne's list of approved devices.
Evaluate ARCserve hardware and software requirements	Page 2-3
Install ARCserve	Page 2-4
Start ARCserve Manager	Page 2-14
Insert a tape into the tape drive	
Configure your tape groups -If you have more than one tape drive, you can separate your drives into different groups to allow parallel streaming. -If you have two or more of the same type of tape drive, you can include them in the same group to allow automated tape spanning.	Chapter 9, "Managing Tapes & Tape Devices"
Back up your files -Open the Backup Manager. -Select an ARCserve Server if you are running ARCserve Manager remotely. -Select your source (the files you want to back up). -Select your destination (the tape you inserted into the tape drive). -Run the job.	Page 2-15 (for the tutorial) See Chapter 4, "Backing Up to Tape" for detailed information.

Hardware and software requirements

To install and use ARCserve, you will need the following hardware and software:

Machine Type	80386 DX or higher computer
Operating System	Windows NT version 3.5 or higher
System Memory	A minimum of 16 megabytes
Disk Space	15-20 megabytes (10 MB for ARCserve files and 5-10 MB for database when jobs are running)

Installing ARCserve

ARCserve has the following installation options:

- Full Setup - install all components of ARCserve on a single machine.
- Custom Setup - install one or more components of ARCserve on a single machine.
- Uninstall - uninstall ARCserve for Windows NT.

Refer to the appropriate section to use the option of your choice.



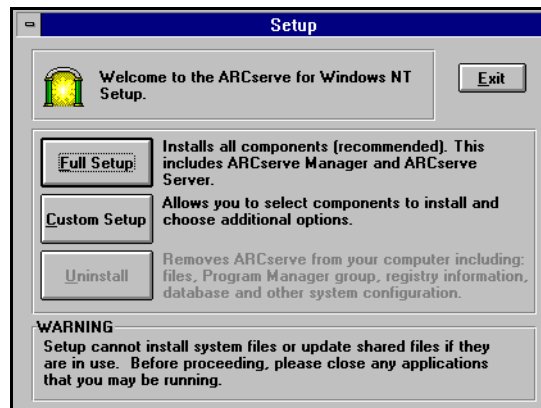
Note that a new user is automatically added to your Windows NT server during installation. Refer to the 'The CheySystem User Account' section on page 2-13.

Full setup

Follow these instructions to install ARCserve on a Windows NT machine:

1. Log on to Windows NT as an Administrator or as a user in the Administrator group.
2. Insert the ARCserve for Windows NT Installation Disk 1 into a local floppy drive.
3. Run setup.exe from the installation disk.


The Setup dialog box opens:



Choose one of the following options:

- **Full Setup** - all components of ARCserve are installed. Go to the next step to continue.
 - **Custom Setup** - select the components of ARCserve that you wish to install. Refer to the 'Custom Setup' section on page 2-10.
4. For a full setup, click the Full Setup button.

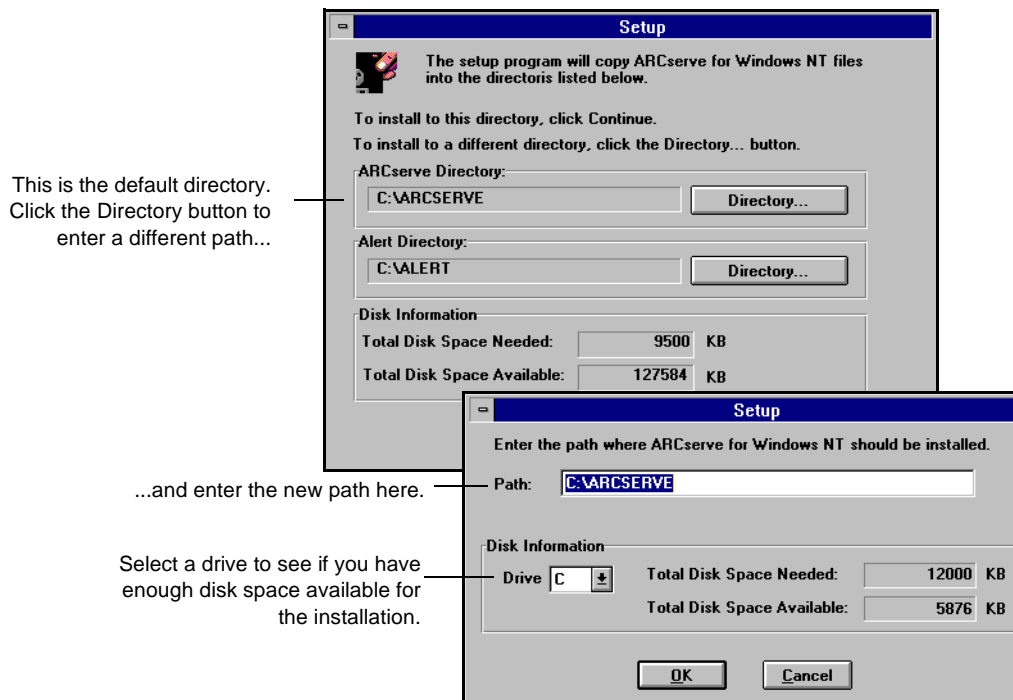
The following dialog box opens:



Enter your name and company name. Verify the information when prompted.

5. Click Continue.

The following dialog box opens:



6. Select your installation directories.

Specify the directories where you wish to install ARCserve and Alert. To change a default directory, click the Directory button next to the directory you wish to change.

Follow these guidelines when choosing a directory:

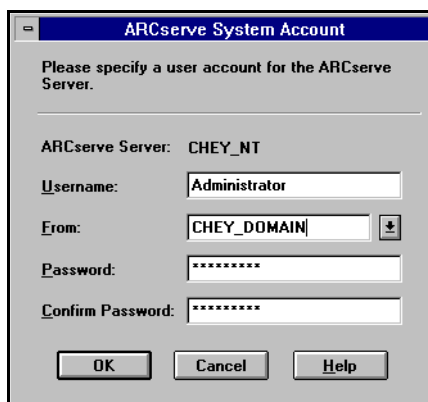
- The directory path you enter must be in the 8.3 (xxxxxxx.yyy) format.
- ARCserve and Alert files must be installed in separate directories on a local drive. Don't install files on a network drive or to a root directory.
- The directory path can be any File Allocation Table (FAT), High Performance File System (HPFS), or Windows NT File System (NTFS) drive. For security reasons, however, we recommend you select an NTFS drive.
- If the specified directory does not exist, it will be created for you.



7. Click Continue.

The ARCserve system is installed and all files are copied to the specified directories.

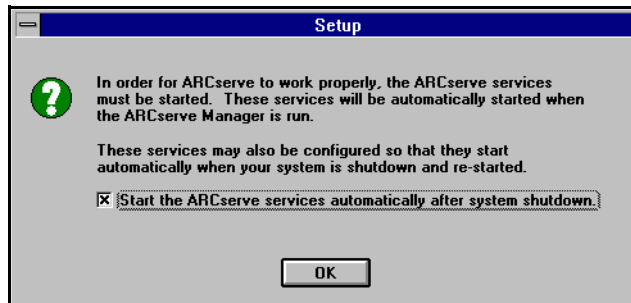
Next, the ARCserve System Account dialog box opens:



Enter a valid username and password for the Windows NT machine on which the ARCserve Server is installed.

8. Click OK.

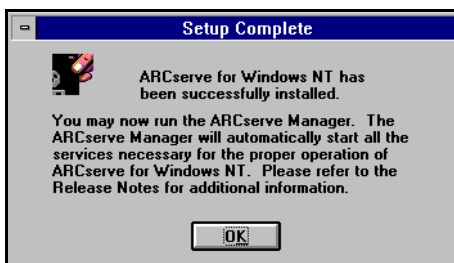
The following dialog box opens:



Select or clear the available option.

9. Click OK.

The following message is displayed.



10. Click OK.

Your installation is complete.

11. Read the important information in the following sections before you run ARCserve:

- 'Using tape devices that were installed through Windows NT Setup' section on page 2-12.

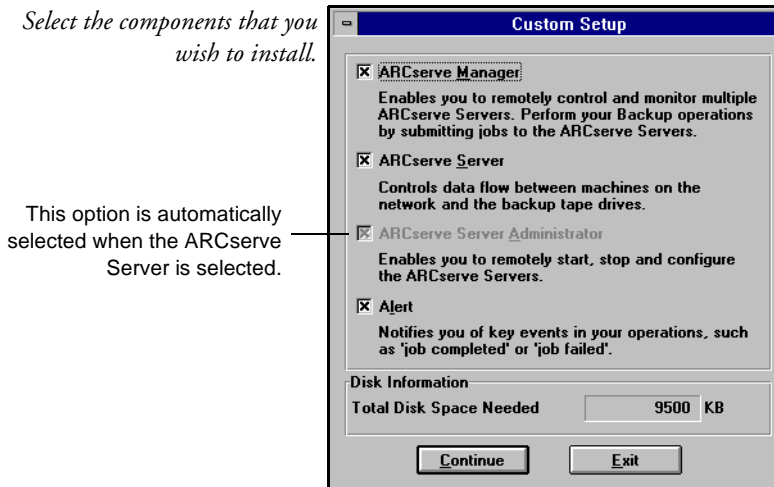
- 'The CheySystem User Account' section on page 2-13.

Custom Setup

To install only selected components of ARCserve Server on a machine, follow these instructions:

1. Follow steps 1-3 in the 'Full setup' section on page 2-5.
2. From the Setup dialog box, click the Custom Setup button.

The Custom Setup dialog box opens:



3. Select the components that you wish to install.
4. Click the Continue button.
5. Follow steps 6-9 in the 'Full setup' section on page 2-5.

Uninstalling ARCserve

The Uninstall option is enabled after you have installed ARCserve for Windows NT. The Uninstall option allows you to remove ARCserve for Windows NT and/or Cheyenne Alert. The Uninstall program shuts down the appropriate services and removes all directory structures and relevant files.

To uninstall ARCserve for Windows NT and/or Cheyenne Alert:

1. From the Windows Program Manger, open the ARCserve for Windows NT Program group.
2. Double-Click the Setup icon.
The Setup dialog box opens.
3. Click the Uninstall button.
You are prompted to select an uninstall option: ARCserve and/or Alert.
4. Choose your uninstall options and click the Continue button.

Uninstall will remove the ARCserve for Windows NT and/or Cheyenne Alert directory structure(s) and all files from your machine.

When ARCserve and/or Alert are removed from your system, the following message box opens:



Using tape devices that were installed through Windows NT Setup

During ARCserve's installation you will see a dialog box that lists any tape devices that were installed through *Windows NT Setup*.

Because they were installed through *Windows NT Setup*, the devices on this list are dedicated to other applications (such as Microsoft Windows NT Backup) and cannot be used by ARCserve.



You will only see this list if you have tape devices that were installed through *Windows NT Setup*.

You will not see this list if you've simply plugged your tape device into its board and then run the ARCserve Setup program.

After installation
you must disable
the device

In order for ARCserve to be able to recognize and use a tape device from this list, you must do one of the following:

- Disable the tape device through *Control Panel - Devices* and then reboot your machine.
- Remove the tape device through *Windows NT Setup - Options - Add/Remove Tape Devices*.

Refer to your Windows NT manuals for more information about using these features.

The CheySystem User Account

When you install ARCserve, the user *CheySystem* is automatically added to your User Account database. This user will be added to the Administrators group and the Backup Operators group. For the Backup Operators group, the following advanced user rights are enabled:

- Act as part of operating system
- Log on as a service

For security reasons, after the installation is completed, you can change the *CheySystem* password (through the *Administrative Tools group - User Manager*). However, if you change the password there, you must also enter the new password under Control Panel's *Services*.



If you do not synchronize the *CheySystem* password in both places, you will not be able to start ARCserve's Job Engine.

Refer to your Windows NT manuals for more information about using these features.

Starting ARCserve - the basic steps

This section explains how to quickly start ARCserve.



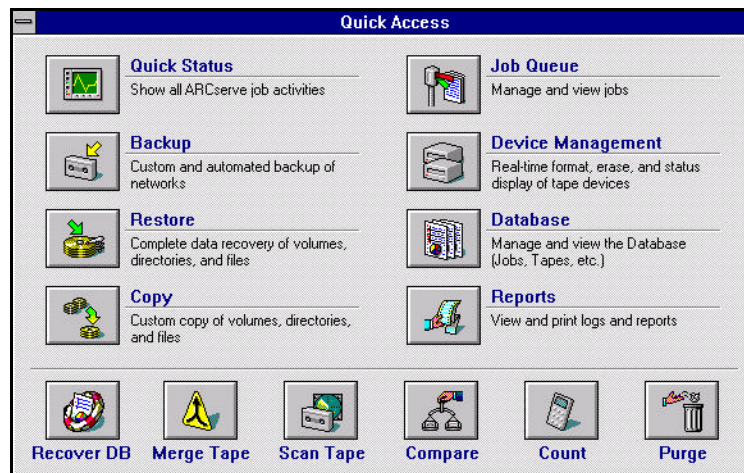
The Job Engine, Tape Engine, and Database Engine (the three components of the ARCserve Server) are automatically started as Windows NT services when you load the ARCserve Manager.

Starting the ARCserve Manager

To start the ARCserve Manager:

1. Double-click the ARCserve Manager icon.

The ARCserve Quick Access window appears. From here you select the ARCserve operation you want to perform.



Backing up your machine - a tutorial

ARCserve has powerful backup capabilities that are easy to use. You can back up shared Windows NT files and directories, as well as entire drives, to tape. In addition, you can back up files stored on NetWare file servers.

For this tutorial we are going to do three things:

- Select a server (to use an ARCserve Server running on a remote machine)
- Format a tape
- Back up your entire machine

The backup we perform will be a very basic backup - that is, we will use ARCserve's default backup options to back up your machine to tape.

After you complete this tutorial and feel comfortable using ARCserve to back up data, you can read through the rest of this manual to discover all of the features and functionality that ARCserve has to offer.

Beginning the tutorial

Follow the steps below to complete the tutorial:

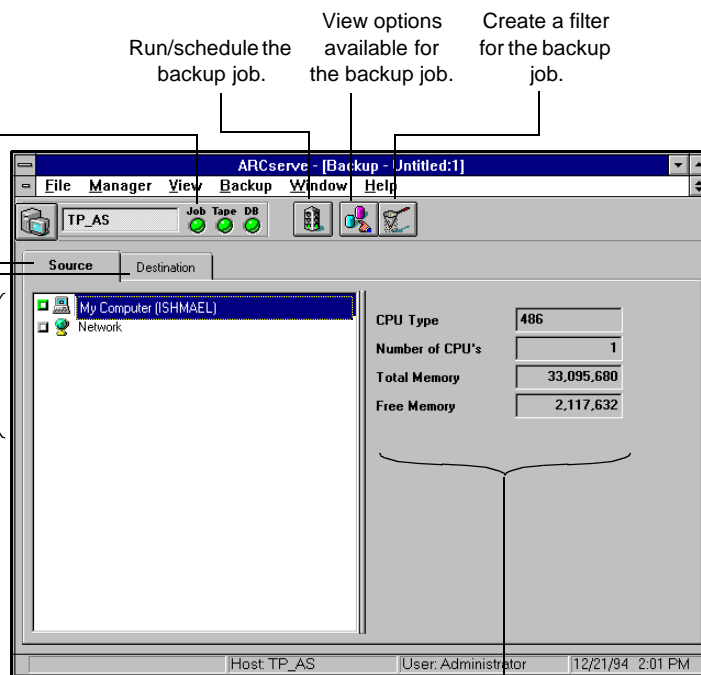


1. Click the Backup button from the Quick Access window .
The Backup Manager appears. The major features of this window are highlighted below:

When these lights are green, ARCserve's engines are running. If a light is red, you can highlight it and click your right mouse button to start the engine.

Use the Source tab to specify the source for the backup job, and the Destination tab to specify the destination for the backup job.

Use the Browser to specify the drives, directories, and files to be backed up. Double-click on the machine name to expand the Browser so you can see the drives on the machine. Then, double-click on a drive to see the directories and files below it.



This pane provides information about the path you have highlighted in the left pane. In this case, the path highlighted is "My Computer", which contains 1 CPU with a total of 32 Mb of RAM, out of which 2 Mb are currently free.

If you wish to select a remote ARCserve Server for this backup job, go to the next step. Otherwise (if you are using the ARCserve Server that is installed on this machine) go to step 8.

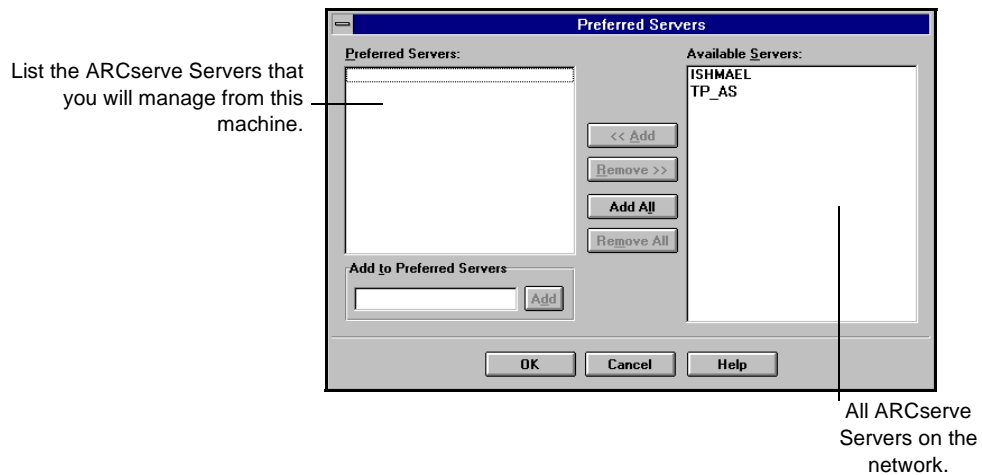
-
- Select a server
2. Click the Select Server button. A small icon representing a server or folder selection, showing a folder with a blue arrow pointing to it.

The following dialog box opens:



3. Click Yes.

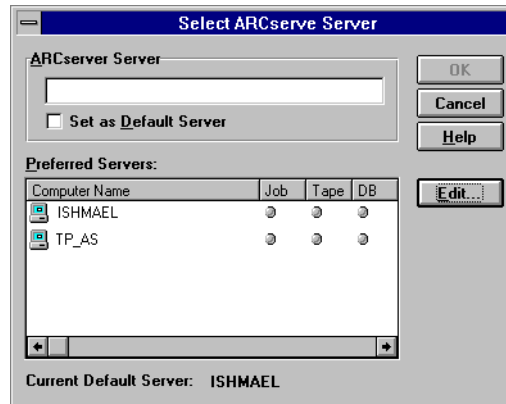
The Preferred Servers dialog box opens:



Create a Preferred Servers list (a list of ARCserve Servers that you will remotely manage from this machine.) Select one or more servers from the Available Servers list, and click <<Add to add them to the Preferred Servers list. Refer to Chapter 3, 'Basic Procedures' for more information about preferred servers.

4. Click OK to continue.

The Select ARCserve Server dialog box opens.



5. Select a server from the Preferred Servers list.
6. Select the option Set as Default Server to make this your default server.
7. Click OK.

A new server is selected for this Backup window.

Select your source - what you want to backup

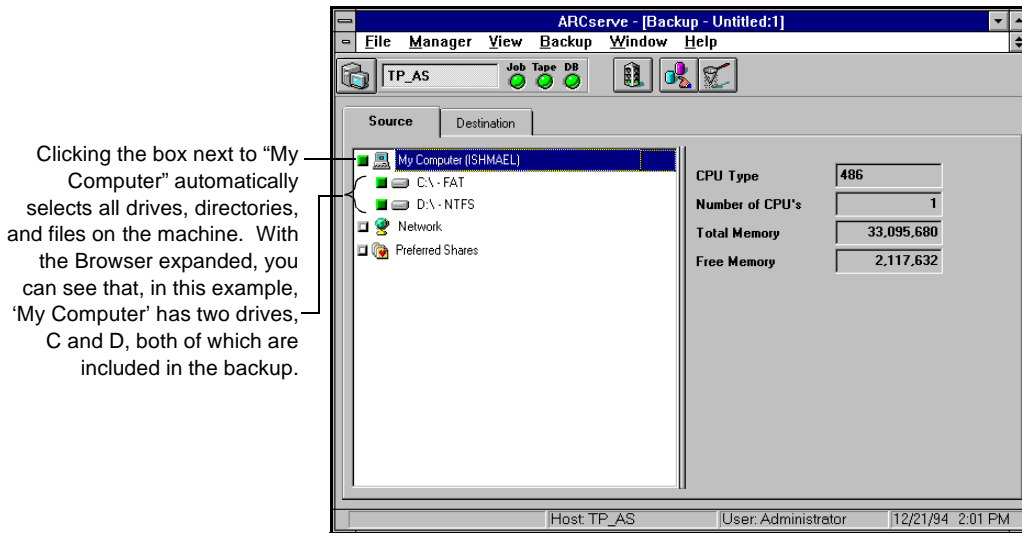
8. From the Backup window, select the Source tab (if it is not already selected).

The Browser displays the machine tree from which you can select machines, volumes (drives), and directories for backup.

9. From the Browser, select "My Computer".

Click the green box next to "My Computer" so that it looks like . This selects all of the drives, directories, and files on your machine.

Your screen should now look similar to the following:



Clicking the box next to "My Computer" automatically selects all drives, directories, and files on the machine. With the Browser expanded, you can see that, in this example, 'My Computer' has two drives, C and D, both of which are included in the backup.

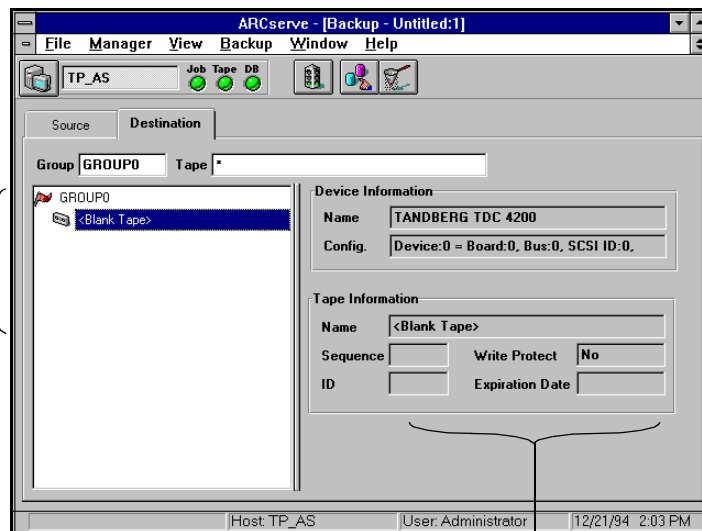


Select your destination - the tape where you want to store the backed up data

10. Select the Destination tab.

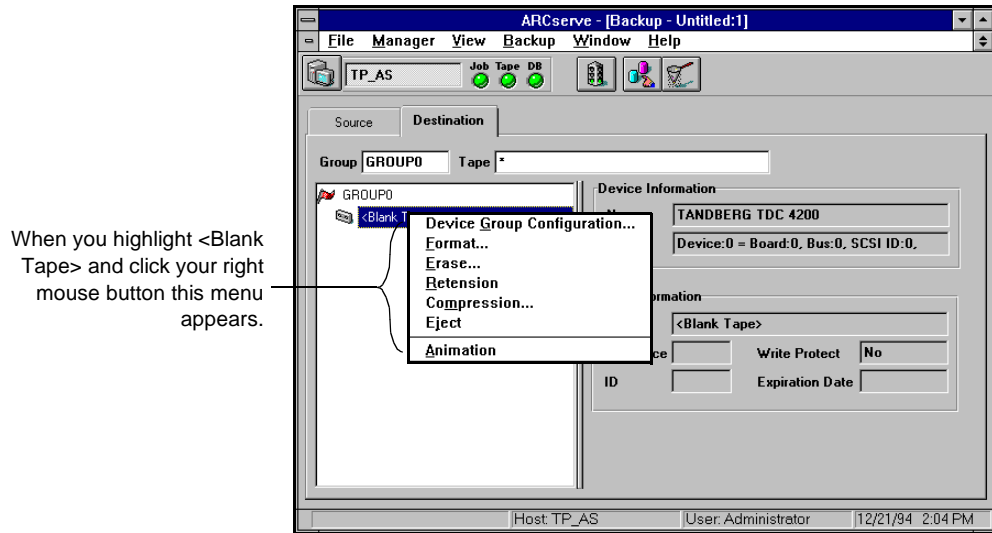
The Browser now displays the tape groups and tapes you can select for your destination.

In the Destination tab the Browser is used to specify which tape or tape group to use for the backup. Highlight the tape or tape group you want to use. When you do, information about the tape and the drive appears in the right pane. (Refer to Refer to Chapter 9, "Managing Tapes & Tape Devices", for information about grouping tape drives.)





This pane provides information about the tape drive and tape you highlighted in the left pane.

-
- Format your tape 11. Highlight <Blank tape> and click your right mouse button .
A menu appears in the Browser:

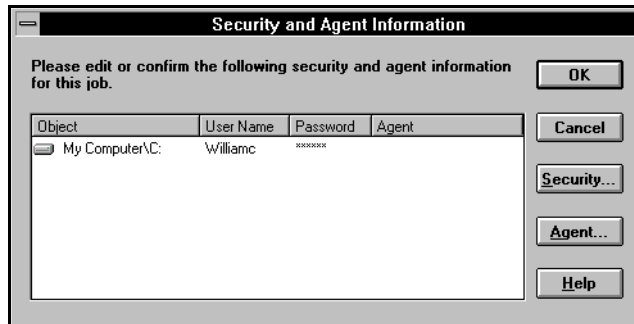


12. Select Format from this menu.
The following dialog box appears:



13. Click the Run/Schedule button. 

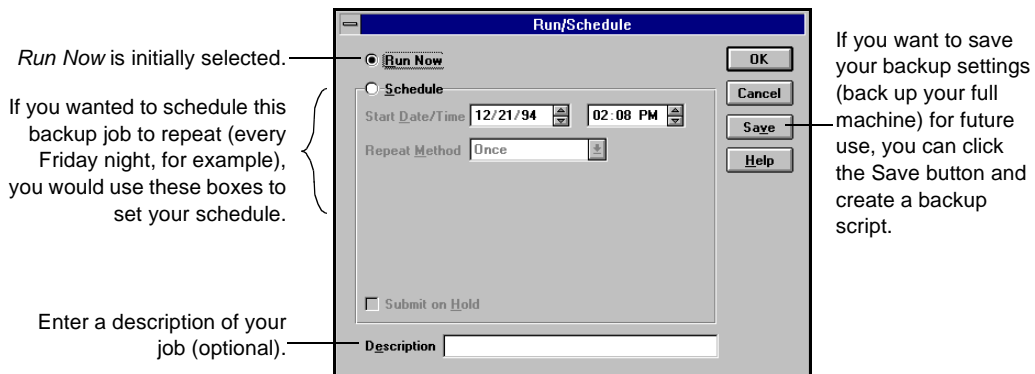
The Security and Agent Information dialog box opens:



This dialog box displays the security and agent information for each machine that is selected as a source. Select the machine “My Computer” and click the Security button to enter the username and password for your machine.

14. Click OK to continue.

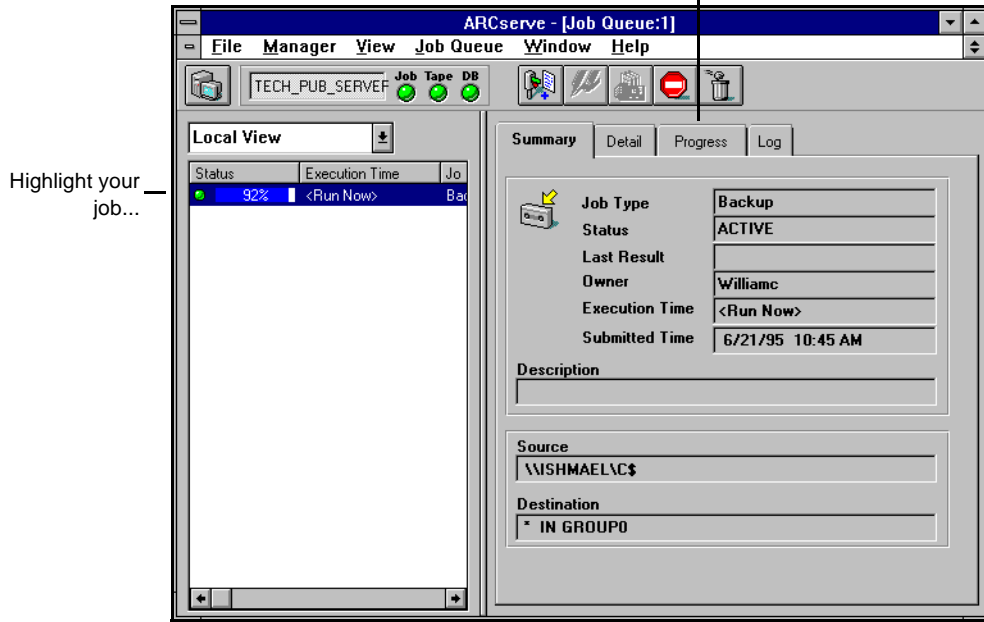
The Run/Schedule dialog box opens:



15. Click OK to run this backup job.

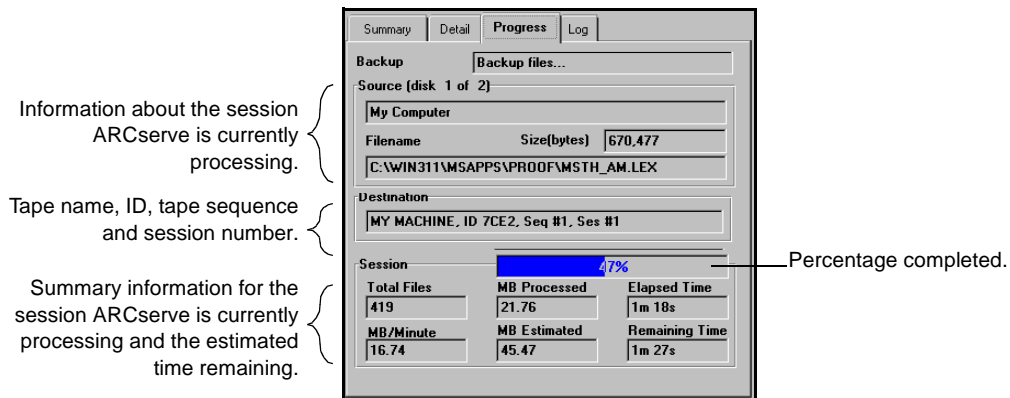
Because you retained the default setting, *Run Now*, the Job Queue Manager is displayed:

...and select the Progress tab to view the progress of your job.



Highlight your job...

16. Highlight your job, and select the Progress tab.
The progress is displayed:



Information about the session ARCserve is currently processing.

Tape name, ID, tape sequence and session number.

Summary information for the session ARCserve is currently processing and the estimated time remaining.

Percentage completed.

3

Chapter

BASIC PROCEDURES

This chapter explains the basic procedures that you must know in order to operate ARCserve for Windows NT.

In this chapter, you will learn:

<i>Page</i>	
3-2	➤ How to start ARCserve
3-5	➤ The basic parts of an ARCserve Manager window
3-7	➤ How to select an ARCserve Server
3-13	➤ How to select machines, files, and directories for ARCserve operations using the machine tree in the Browser
3-26	➤ How to select tape groups and tapes for ARCserve operations using the tape tree in the Browser
3-30	➤ How to access on-screen and online help

Starting the ARCserve Manager

The ARCserve Manager enables you to perform your backup, copy, restore, and related operations.

To start the ARCserve Manager:

1. Double-click the ARCserve Manager icon in the ARCserve for Windows NT program group.

The ARCserve desktop opens and the Quick Access window is displayed.



2. Select an option from the Quick Access window to continue.

Customizing Quick Access

Quick Access can be displayed in one of the following ways:

- As a detailed window
- As a toolbar
- Not at all

This section describes how to set your preference for the way Quick Access is displayed.

Quick Access window

1. From the View menu, select Quick Access Detail.

This is the default view. The Quick Access window is shown below:



Quick Access toolbar

1. From the View menu, select Quick Access Toolbar.

The Quick Access toolbar appears at the top of the ARCserve Manager window, as shown below:




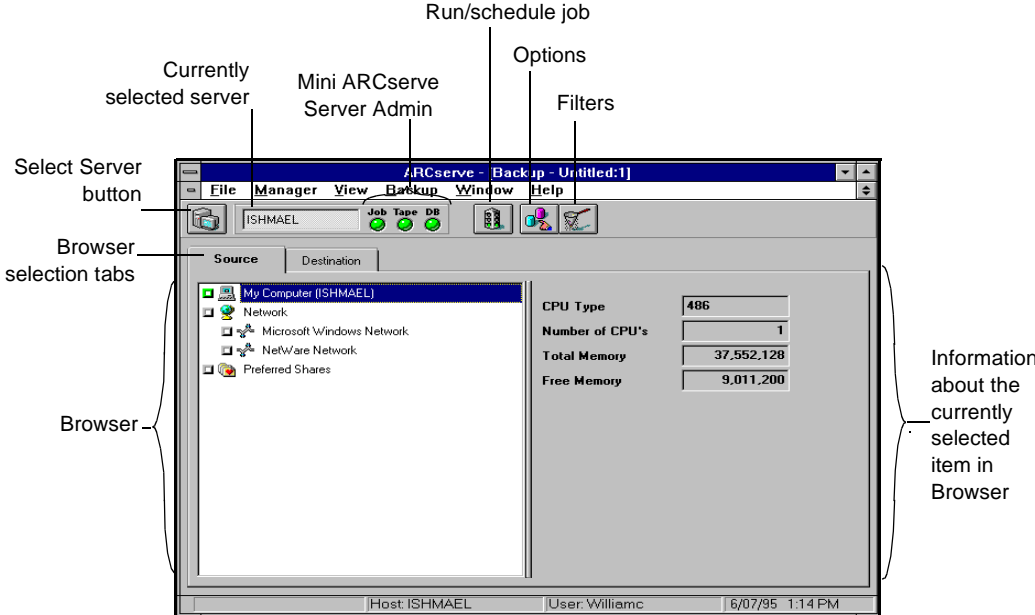
No Quick Access displayed

1. From the View menu, select Quick Access Detail.
This displays the Quick Access window.
2. Close the Quick Access window.
3. From the View menu, select Save Quick Access Settings Now.

From now on (until you select another setting), Quick Access will not appear on your desktop, either as a window or toolbar.

Basic ARCserve Manager features

The following figure illustrates the basic features of an ARCserve Manager window:



Managing your ARCserve Servers

The ARCserve Manager enables you to remotely manage multiple ARCserve servers within your Windows NT domain. You can use the ARCserve Manager to submit job requests to an ARCserve Server, and monitor the status of your requests. You can:


- Open a Manager window, such as the Backup window or the Restore window, for each job request.
- Easily select an ARCserve Server using the Select Server button on the Manager window.
- Manage all backup operations for the selected ARCserve server.
- Open a new window for each additional job you wish to submit to a server. You can select a different ARCserve Server in each open window.
- Switch between job requests on different servers by changing your Window focus.

For example, you can open a Backup window, select server A, and schedule a backup for this server. Now open a Restore window, select server B, and schedule a restore for server B. Now switch to your Backup window, and check on the status of your backup job and server A. Now open a second backup window select server C, and so on.

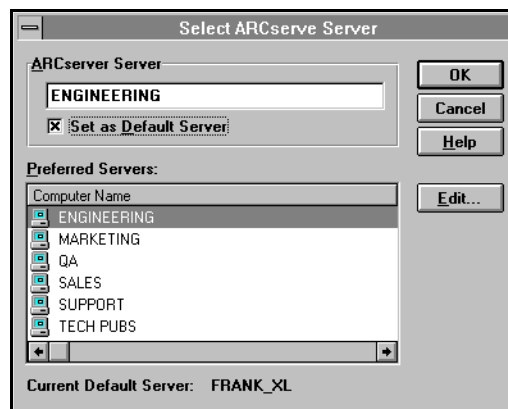
Selecting an ARCserve Server

You can select an ARCserve Server from any Manager window that submits job requests to a server, such as the Backup window or the Restore window. You can select a different server from each open Manager window.

To select an ARCserve Server from a Manager window:

1. Click the Select Server button  on the mini Server Admin.

The Select ARCserve Server dialog box opens.




2. Select a server from the Preferred Servers list.
3. Select the option Set as Default Server to make this your default server.
4. Click OK.

A new server is selected for this Manager window.

Editing the Preferred Servers list

You can create your own Preferred Servers list containing the ARCserve Servers that you will manage from the ARCserve Manager. You can add and remove ARCserve Servers to the Preferred Servers list as needed. To edit the Preferred Servers list:

1. From the Select ARCserve Server dialog box, click Edit. The Edit Preferred List dialog box opens:




2. Add and remove servers to the Preferred List as follows :

Task	Procedure
Add a server	<p>Either:</p> <ul style="list-style-type: none"> • Select the server in the Available Servers list and click <<Add. • Enter the name in the Add to Preferred List box and click Add.
Remove a server	Select the server in the Preferred List, and click Remove>>.
Add all servers	Click Add All.

Task	Procedure
Remove all servers	Click Remove All.

3. Click OK to close the dialog box.
The Select ARCserve Server dialog box is updated.



Controlling the ARCserve Server engines

You can use the mini Server Admin tool to view and change the state of the ARCserve Server engines. For more information about the ARCserve Server engines, refer to Chapter 14, “Administrating the ARCserve Server”.

Determining engine status

The color of the lights in the mini ARCserve Server Admin indicates the current state of each ARCserve engine.

- Green light - the engine is running
- Red light - the engine is stopped
- Yellow light - the engine is paused

Changing the state of an engine

To change the current state of an engine:

1. Place the cursor over the engine’s status color and click the right mouse button.

The Mini ARCserve Server Admin menu appears.



These three menu options change the status for the individual engine only.

These two menu options change the status of all the engines at once.

You can also start and stop the ARCserve engines by selecting the appropriate option from the File menu.

2. Choose to start, stop, or pause the engine or all the engines.

The engine’s status color will change to reflect the engine’s new state.

Starting the engines as needed

ARCserve Manager can automatically start the ARCserve engines *as needed* during your operations. To start the ARCserve engines as they are needed, select the option *Auto Start Engines* from the File menu.

For example, if the *Auto Start Engines* option is selected, then the *Tape engine* is automatically started when you open the *Backup window*. Similarly, the *Tape and Database engines* are started when you open the *Restore Window*.

Starting the engines on Windows NT startup

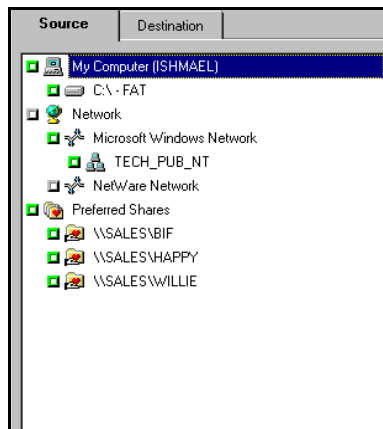
The Setup program provides an option to automatically start your engines on Windows NT startup. You can also specify this option by configuring the ARCserve engine services in the Windows NT Services program. (Click the Services icon in the Windows NT Control Panel program group to access this program.) Refer to your Windows NT manuals for more information.

Using the Browser

The Browser facility is used in many Manager windows to enable you to choose the source and destination for your backup, copy, and restore operations. The Browser operates in two modes:

- **Machine tree:** view and select files, directories, drives, shares, machines, volumes, and servers.
- **Tape tree:** view and select tape groups and tapes.


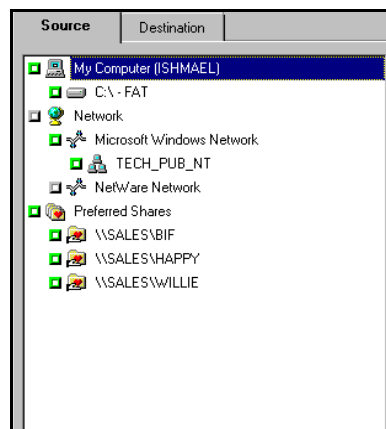
The Browser mode is determined by the type of operation (e.g. backup or restore) and whether the *Source* or *Destination* tabs are selected. In the following figure, a machine tree is shown in the Browser:



Using the machine tree in the Browser

The machine tree in the Browser enables you to view and select files, directories, drives, shares, machines, volumes, and servers for your backup, copy, restore, and related operations.







A machine tree is shown in the following figure:



Basic elements
of the machine
tree


The basic elements of the machine tree are described in the following table:

Symbol	Represents
	Available network on your LAN
	Network Provider (e.g. Microsoft Windows network or NetWare network)
	Windows NT Domain or Workgroup
	Windows NT machine

Symbol	Represents
	Novell NetWare file server
	Preferred Shares
	Disk drive
	Directory
	File
	Indicates that the file, directory, or volume is compressed

Selecting an item in the machine tree

To select an item in the machine tree, click the green selection box to the left of the item. A solid box indicates that the item has been selected. A partially filled box indicates that some files or directories below this level have been selected.



Selecting your local machine

The Browser uses the following label to indicate the local machine on which the ARCserve Manager is running:

My Computer (machine name)

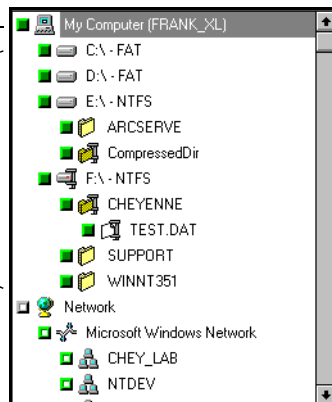
To select your local machine:

1. In the Browser, select *My Computer*.

Click the green selection box next to *My Computer*; all volumes or drives, directories, and files on your machine are selected.

Clicking the green box next to *My Computer*...

...automatically selects everything on the machine, as well.



Selecting a Windows NT domain or workgroup

Selecting a Windows NT domain or workgroup automatically selects each Windows NT machine that belongs to it.



1. In the Browser, expand the tree under Network.

Double-click Network to expand the tree. The types of available networks are listed.



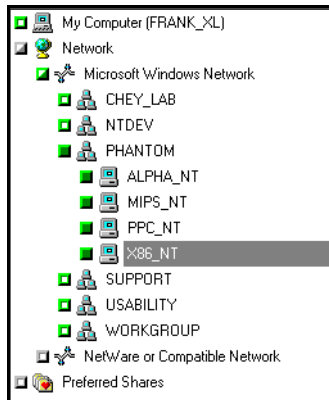
2. Expand the tree under Microsoft Windows Network.

A list of Windows NT domains and workgroups on your network appears.



3. Select the domain or workgroup.

When you select a domain or workgroup, all the machines that belong to it are selected.



4. (Optional) Provide security information for the domain.

Refer to the 'Providing security information for remote machines' section on page 3-25 for details.

Selecting individual Windows NT machines



1. In the Browser, expand the tree under Network.
Double-click Network to expand the tree. The types of available networks are listed.




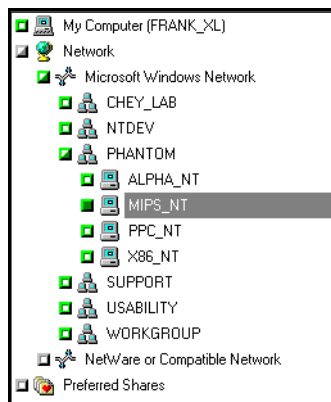
2. Expand the tree under Microsoft Windows Network.
A list of Windows NT domains and workgroups on your network appears.



3. Expand the tree under the domain or workgroup that the machine belongs to.
A list of Windows NT machines appears.

4. Select the machine.

To select a specific Windows NT machine () , expand the tree under the domain or workgroup that it belongs to, and then select the machine.



5. (Optional) Provide security information for the machine.
Refer to the 'Providing security information for remote machines' section on page 3-25 for details.

Selecting a NetWare server



For the most complete backup of NetWare servers, use the optional Client Agent for NetWare. See chapter 16, “About the ARCserve Agents” for more information.

To select a NetWare server:



1. In the Browser, expand the tree under Network .

Double-click Network to expand the tree. The types of available networks are listed.

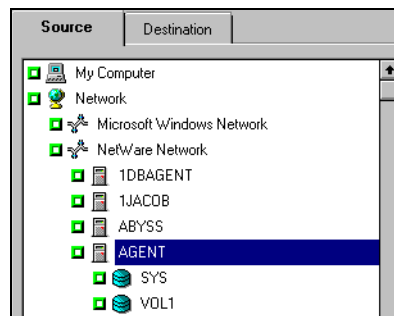


2. Expand the tree under NetWare Network .

A list of NetWare file servers on your network appears.



3. Select the file server .



4. Provide security information for the domain .

Refer to the ‘Providing security information for remote machines’ section on page 3-25 for details.

Selecting an entire drive or directory

To select an entire drive or directory:


1. In the Browser, double-click on the name of the machine that has the drive you want.

The Browser now shows you a graphical tree of the drives available on the machine.



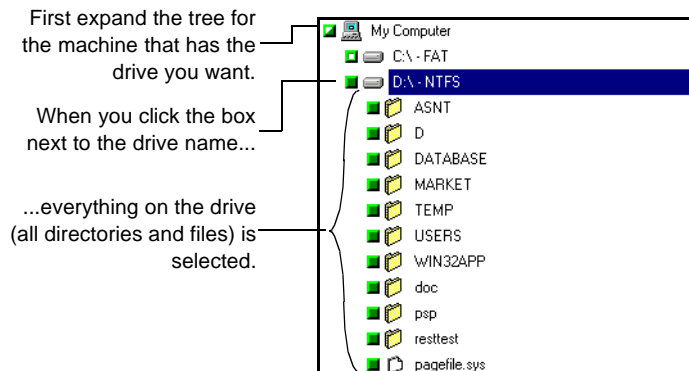
For a remote machine, you may need to provide security information before you see the graphical tree. Refer to the 'Providing security information for remote machines' section on page 3-25 for details.


3

Selecting an entire drive 

2. Select the drive.

If a green box appears next to the drive, click it so that it is completely filled. This automatically selects the entire drive for the operation. All directories and files on the drive are selected.



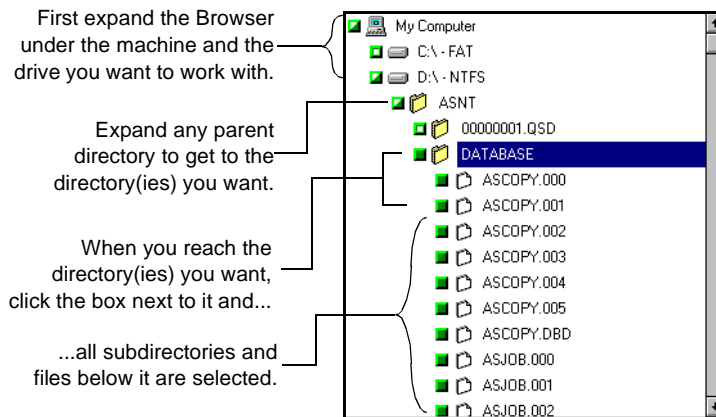
Selecting a directory 

To select a directory, double-click on the drive instead of simply selecting it. Double-clicking on the drive will display a list of directories on that drive.

3. Select the directory.

If a green box appears next to the directory, click the box so that it is completely filled.

If the directory you want is a subdirectory, continually expand the directory tree from its parent directories until you see the directory you want.



Adding Preferred Shares to the machine tree

The Preferred Shares are a collection of your favorite backup shares. (A share is a shared drive or directory.) You can manually add individual share points to the Preferred Shares tree; the share point is remembered and displayed regardless of the status of the network connection. This provides a quick access to commonly used shares on your machines.

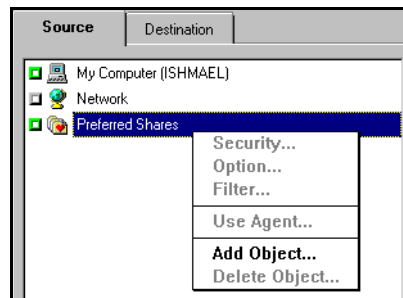
This feature can be used to back up share points provided by other operating systems that are capable of creating “shares” in a Windows NT network, such as UNIX NFS shares or OS/2 LAN Server shares.

This feature is also useful for users that frequently back up “shares” within a network, instead of “machines”. Shares are considered “named” short-cuts or aliases to machine drives and directories on the network. Shares perform the same function as drive mappings, with the advantage of allowing the user to provide a descriptive name to the share.

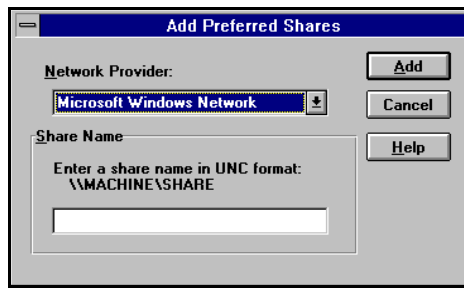
To add a Preferred Share to the machine tree:

1. In the Browser, click Preferred Shares with your right mouse button.

A pop-up menu is displayed:



-
2. Choose Add Object... from the popup menu.
The Add Preferred Shares dialog box opens:



3. Select a Network Provider.
4. Enter a share name in Uniform Naming Convention (UNC) format (e.g. \\MACHINE\SHARE).
5. Click Add.
The machine is added to the Preferred Shares tree.
6. Repeat steps 3-5 for each machine that you wish to add.
7. Choose Close from the Control Menu when done.

Adding objects to the machine tree

You can add objects to the machine tree that are unavailable when you set up a job. For example, you can add a machine that is temporarily unavailable and schedule a backup/restore for the machine.

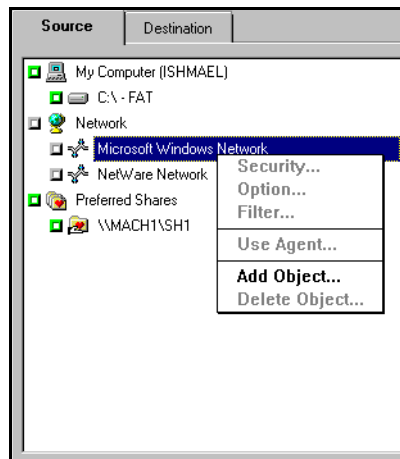
You can add the following types of objects to the machine tree:

- Windows NT domains (or workgroups), machines, or shares
- NetWare servers and volumes

For example, to add a Windows NT domain follow these steps:


1. In the Browser, click Microsoft Windows Network with your right mouse button.

A pop-up menu is displayed:



2. Choose Add Object from the pop-up menu.

The Add Object dialog box opens:



3. Enter the Domain name and click Add .

A new Windows NT domain is added to this node of the machine tree.

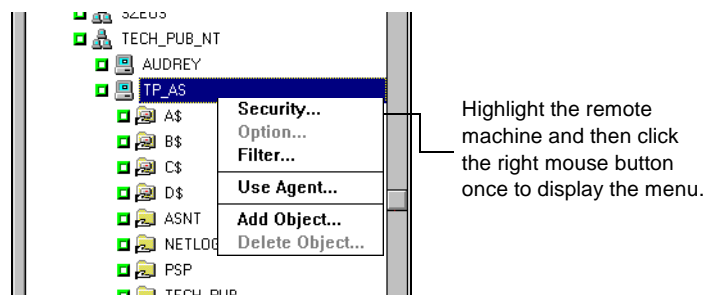
Providing security information for remote machines

You must provide a valid username and password in order to browse a remote machine. By default, ARCserve will use the username and password provided for the local machine to access the remote machine. If this information is not valid, you will be prompted to enter valid information when you try to expand the node for the local machine.

You can also enter a username and password for a remote machine as follows:

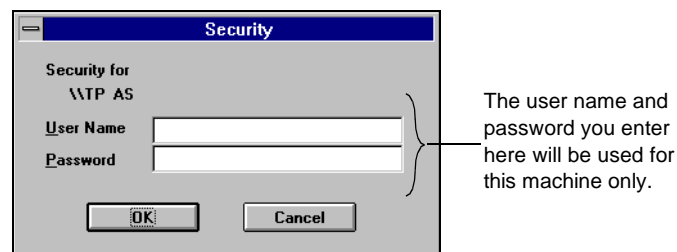
1. In the Browser, click the remote machine with your right mouse button.

A pop-up menu is displayed:



2. From the pop-up menu that appears, select Security... and enter the information.

The Security dialog box opens.



Using the tape tree in the Browser


The Browser is also used to select the tape groups and tapes that you want to use for your jobs. Depending on the job you want to perform, you may need to select a tape or tape group as your source or destination. Whether the job's source or destination, the method for selecting a tape or tape group is the same.


Basic elements of the tape tree in the Browser


The following screen shows the basic elements of the Browser when selecting tapes or tape groups:

Highest level of the tree, this is the name of the tape group.

Tape in one of the devices in the tape group.



 Represents a tape group.

 Represents a tape.

Selecting a tape group only

Follow the directions below to select a specific tape group.

1. Highlight the name of the tape group you want to use.
Click once on the name of the tape group so that it is highlighted. The name of the tape group appears in the *Group* field.



An asterisk (*) appears in the *Tape* field. This means that any tape will be used for the job.


If you select a tape group only, the first tape found in the first device in the tape group will be used for the operation. The job will be appended to this tape.

Selecting a tape

1. Expand the tree for the tape group you want to use.
The Browser shows you all the tapes found in devices assigned to the tape group. The name of this tape group appears in the *Group* field.

The name of the tape group you double-clicked in the Browser area appears here.

The Browser lists all the tapes found in this tape group's devices.



2. Highlight or specify the name of the tape you want to use.

Appending the job to any tape

By default, an asterisk (*) appears in the *Tape* field. If you leave the asterisk in this field and make sure the tape option, *Append to Tape* is selected.

Appending the job to a specific tape


If you want to append the job to a specific tape, enter the name of the tape in the *Tape* field and make sure the *Append to Tape* tape option is selected.

Reformatting the tape with a specific name

If you enter a name in the *Tape* field (instead of selecting an existing tape) and select an "Overwrite..." tape option, ARCserve will reformat the tape in the drive with the name that you specified.


Reformatting the tape without a specific name

If you leave the asterisk (*) in the *Tape* field (instead of entering a name) and select an "Overwrite..." tape option,



ARCserve will reformat the tape in the drive, using the current date and time for the tape name.

After you highlight the tape in the Browser, it appears in this field.



Using ARCserve's menu bar

ARCserve's menu bar appears at the top of every window. This menu bar contains drop-down menus by which you can select a function to be performed.

Whenever the step-by-step instructions in this manual tell you to select a function from a menu, use your mouse to highlight the menu option.

The underlined letters are the key to keyboard shortcuts

You can also use keystrokes to choose menu options. From the Main menu, or the Quick Access screen, you can press the ALT key together with the underlined letter to select a menu.

Once the menu is displayed, you can select an option from it by pressing the underlined letter in the option name.

Getting help

There are two types of help available in ARCserve:

- On-screen descriptions of graphical command buttons
- Online help system

On-screen descriptions of buttons

Some ARCserve command buttons have words on them, such as “Query” and “Version History”, while other command buttons have graphics that represent the function of the button. While these graphics are meant to be easily understood, on-screen descriptions are also provided, which tell you the function of each button.

To access on-screen descriptions of graphical buttons:

1. Place the tip of the cursor over the button for a few seconds.

A brief description of the button appears just below the cursor and at the bottom of the screen.

Online help system - provides instructional help and information on ARCserve error messages

The online help system provides instructional and general information about ARCserve. It also provides information on the ARCserve error messages.

The information provided by the online help follows a path similar to the format used in this manual. You select a concept from a list of topics (similar to a Table of Contents) and then general information about that topic is displayed (similar to reading the overview of a chapter in this manual). From here, you can display more detailed information, including specific instructions (similar to turning to a specific page or section in a chapter). Or, if an ARCserve dialog box or screen has a Help button (a button with the word “Help” on it), you can click it to see information specific to that dialog box or screen.

To access online help, follow the directions below:

1. From the Help menu, select Contents.

The Online Help system is started, and a list of topics available for help is displayed.

Other sources for help

As you use ARCserve, you may find that you need additional information that may not be provided in this manual. Following is a list of other sources for help in using ARCserve:

- Microsoft Windows NT documentation - for any questions you might have about using the Windows NT interface.
- Windows NT Resource Kit - for valuable information about Windows NT not provided by the Windows NT documentation. Information and help about Windows NT error messages are provided, also.



4

Chapter

BACKING UP TO TAPE

Use the ARCserve Backup Manager to back up Windows NT machines and shared network drives. This is an important step in protecting your valuable data.

In this chapter, you will learn:

Page

- | | | |
|------|---|---|
| 4-4 | ➤ | How to use the Backup Manager |
| 4-5 | ➤ | How to perform a quick, basic backup |
| 4-10 | ➤ | How to use advanced backup features and options |

Backup basics

ARCserve's powerful backup capabilities are easy to use. With them, you can back up:

- Your local Windows NT machine
- Remote Windows NT machines on your Microsoft network
- NetWare file servers (Use the optional Client Agent for NetWare for complete backup.)

Backing up remote Windows NT machines

ARCserve for Windows NT allows you to back up any Windows NT machine on your Microsoft network. Backing up a Windows NT machine can be done one of two ways:

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

Administrative shared drives

By default, Windows NT creates an administrative share for the root directory of each drive on the machine. Unless the administrative share is deleted, you can select it for backup and ARCserve is able to back up all the files and directories on the drive.

User shared drives and directories

In addition to creating default administrative shares, Windows NT also allows you to designate specific files and directories as shareable. The owner of these files and directories has specified that they are accessible to all. These files and directories are displayed in ARCserve's Browser, and with a click of the mouse button, you can quickly select these specially shared files and directories for backup.

Perform domain- or workgroup- level backups

Do you want to quickly back up all the machines belonging to a specific domain or workgroup? Since ARCserve separates and lists Windows NT machines by the domain or workgroup to which they belong, this is easily done by selecting the name of the domain or workgroup. Doing this automatically selects all the machines that are members of the domain or workgroup for backup.

Get familiar with the basics, then use advanced features and options

Once you get accustomed to the basics of setting up a backup job, explore the backup options, filtering, and scheduling techniques to really take advantage of ARCserve's power and flexibility.

Backup options are specifications that you can make or change to enhance and define each job. To change the default backup options, refer to the 'Advanced backup features and options' section on page 4-10.

Filtering allows you to be even more explicit about the files and directories that you want backed up. Suppose you want to back up the majority of files and directories on a particular Windows NT drive. Instead of using the Browser to select each directory and file that you want to back up, use the Browser to select the entire drive, then add filtering specifications to exclude the directories and files that you don't want included in the job. For details about filters and how to use them, refer to Chapter 7, "Customizing Your Jobs".

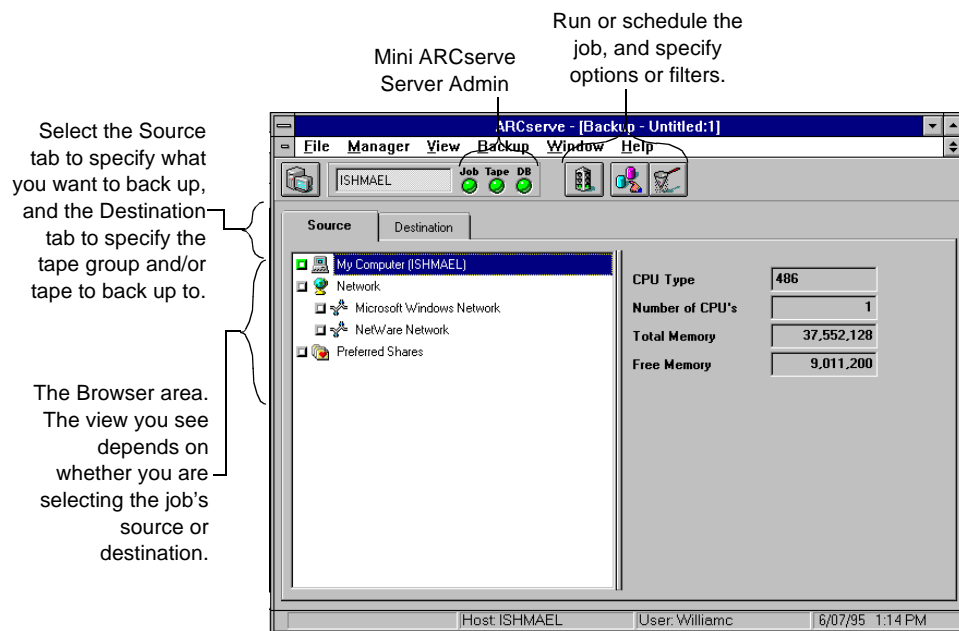
Another way of enhancing your backup scheme is to apply different scheduling techniques to your jobs. Instead of constantly re-submitting jobs that you want to run frequently, apply a repeat method to these jobs and let ARCserve handle the resubmission for you. For details about scheduling your job, refer to Chapter 7, "Customizing Your Jobs".

Save important jobs to a script file

Make it a habit to frequently save backup job settings to a script file. It will make it easier for you to resubmit jobs to the queue if and when you need to. For more information about scripts, refer to Chapter 7, "Customizing Your Jobs".

Backup Manager

The Backup Manager is shown below:




Backing up files - the basic steps

Following are the basic steps for making a backup - that is, using the default backup options to back up your data to tape. Specifying different options for a backup is covered later in this chapter.

Sean wants to back up some important information on his network. Since he has one tape group configured for his machine, he will simply append this job to one of his tapes, SEAN_BACKUP.

Basic steps for making a backup



Using the above example, follow these steps to quickly back up data:

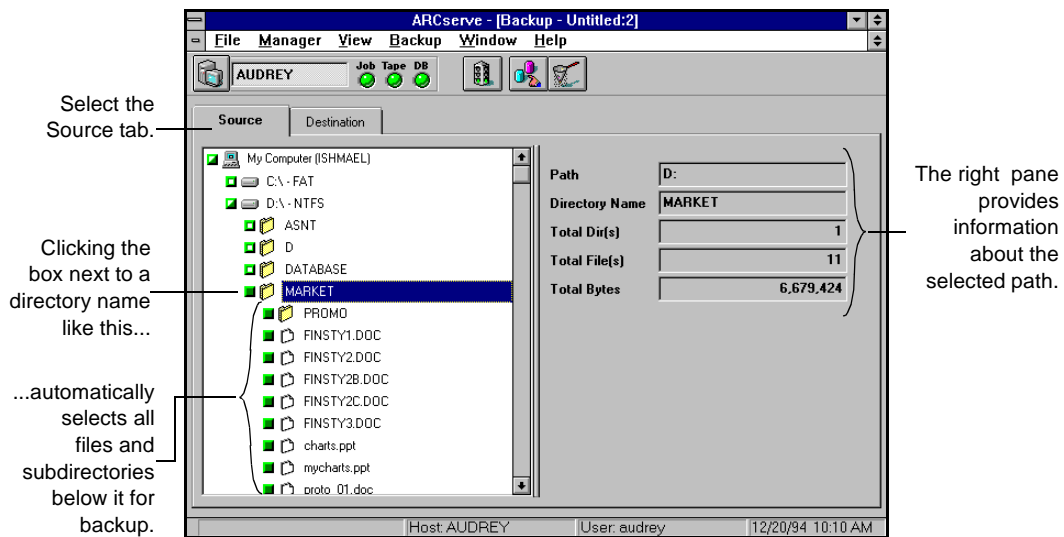
1. Click the Backup button from the Quick Access window .
The Backup Manager appears.

2. Select what you want to back up .

Select the Source tab in the left pane of the Backup Manager window. The Browser displays the machine tree, which shows available networks and machines. Select the items from the machine tree that you wish to back up.

Refer to Chapter 3, 'Basic Procedures', to learn how to select items from the Browser.

Sean selects the items from the machine tree that he wants to backup. He selects one Windows NT domain, two NetWare file servers and the "MARKET" directory on his machine. (To select the "MARKET" directory, he expands the tree under "My Computer", as shown in the following diagram.)



Refer to Chapter 3, "Basic Procedures", for details on using the tape tree in the Browser

- Specify the tape group (and tape, if you have a specific one in mind) you want the files backed up to .

Select the Destination tab to display the tape tree in the Browser. Then, use the Browser to select the tape group and tape you want to use for the backup job.

Sean's tape group is called GROUPO. He places the tape, SEAN_BACKUP, in a device, and proceeds to select GROUPO and SEAN_BACKUP as the job's destination.

The screenshot shows the ARCserve Backup software interface. The window title is "ARCserve - [Backup - Untitled:2]". The menu bar includes "File", "Manager", "View", "Backup", "Window", and "Help". The "Destination" tab is selected, showing a tree view on the left with "GROUPO" expanded to show "SEAN_BACKUP". The right pane displays "Device Information" and "Tape Information".

Device Information:

Name	HP HP35480A
Config.	Device:0 = Board:0, Bus:0, SCSI ID:1.

Tape Information:

Name	SEAN_BACKUP	
Sequence	1	Write Protect
ID	B842	Expiration Date
		12/31/95 1

Annotations on the left side of the screenshot:

- "Select the Destination tab..." points to the "Destination" tab.
- "...then use the Browser area to select the tape group and tape you want to back up to." points to the tree view.

Annotations on the right side of the screenshot:

- "What you select in the left pane is placed in the fields here. Or, you can manually enter the information instead." points to the "Group" and "Tape" fields.
- "Details about the tape group or the tape (whichever is highlighted at the time) is displayed here." points to the "Device Information" and "Tape Information" panes.

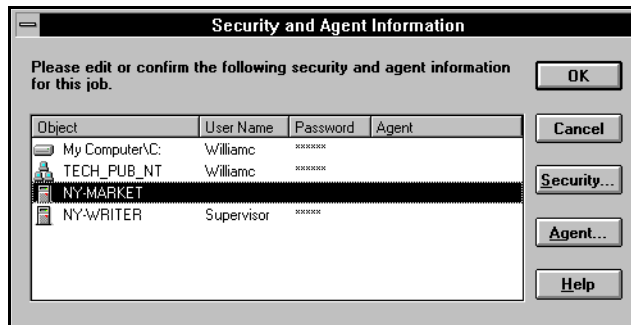
At the bottom of the window, it shows "Host: AUDREY", "User: audrey", and "12/20/94 10:15 AM".



-
4. Click the Run/Schedule button.



The Security and Agent Information dialog box opens:

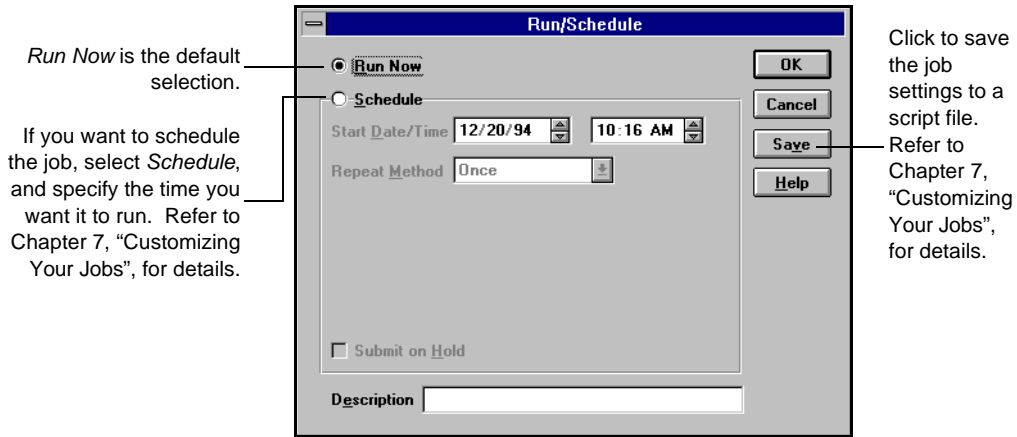


This dialog box displays the security and agent information for each machine and/or share that is selected as a source. To change this information, select the machine, then select either the Security button or the Agent button.

Sean enters the username and password for each machine that he has selected as his source. He is not using an agent, so he leaves the agent information blank.

5. Click OK to continue.

The Run/Schedule dialog box opens.



The Run/Schedule dialog box enables you to run the job, or schedule the job for a later time.

Since Sean wants to back up his files now, he clicks the Run/Schedule button and selects Run Now to execute the job immediately.

To save this job to a script file, click the Save button in this dialog box.

6. Click OK to submit this job.



Advanced backup features and options

Now that you know how to perform a basic backup, you are ready for ARCserve's advanced backup features. These include using hidden, right-click menus provided by the Backup Manager and specifying job options.

Hidden menus in the Browser

ARCserve provides hidden menus in the Browser to make specifying your backup job easier. These menus are accessed by clicking the right mouse button after highlighting a machine, tape group, or tape. The menu that you see depends on whether you are selecting your source or your destination.

Source Context menu

If you are selecting the job's source, the hidden menu allows you to specify local filters and options (filters and options that apply only to the machine you have highlighted), as well as provide a user login name and password to access the machine.



For details about specifying **local security information**, refer to Chapter 3, "Basic Procedures".

For details about specifying **local options**, refer to the 'Applying local filters and options' section on page 4-12.

For details about specifying **local filters**, refer to Chapter 7, "Customizing Your Jobs".

Destination Context menu

If you are selecting the job destination, the hidden menu allows you to configure tape groups or perform maintenance functions on the tape.



For details about configuring tape groups and performing maintenance functions on tapes, refer to Chapter 9 , “Managing Tapes & Tape Devices”.

Applying global options to your job

The backup options you specify in the Options dialog box ensure that you get the most reliable and efficient backups possible. Whereas the local options apply only to the highlighted node or node drive, global options apply to the entire job. For details on applying global options, refer to the ‘Providing security information’ section on page 4-14.

Applying local filters and options

ARCserve provides a hidden menu that allow you to specify information that should be applied only to the drive that you have highlighted.




For detailed information on filters and how to apply them locally to a specific node, refer to Chapter 7, “Customizing Your Jobs”.

To apply a local filter or option, follow the steps below:

1. Select the Source tab.
2. Highlight the node or the drive on the node and press the right mouse button.

The Source Context menu appears:

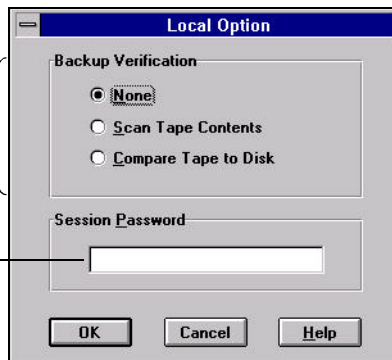


Specifying local options

3. Choose Option... from the Source Context menu.
The Local Option dialog box appears:

Data backed up from the drive will be verified according to your selection.

The password you provide here will be required to restore the data that was backed up from this drive.



Any options you specify in this dialog box will be applied ONLY to the drive you highlighted.



For descriptions of the Verification options, refer to the 'Verifying your backup' section on page 4-28.



Providing security information

You must provide a valid username and password for the machines or shares that you wish to back up. You can enter the username and password:


- By clicking on the machine or share with your *right* mouse button.
- When attempting to expand the machine or share to view directories. (You will be prompted for the information.)
- When running a backup job. (You will be prompted for the information.)

Browsing

The Browser uses your local username and password by default in order to access machines or shares for browsing. If your local username and password are not valid on the remote machine or share, then you will be prompted to enter the information when you attempt to expand that node.

You can also enter a username and password from the Browser as follows:

1. Select the machine with your right mouse button.
A popup menu is displayed.
2. Select Security from the pop-up menu.

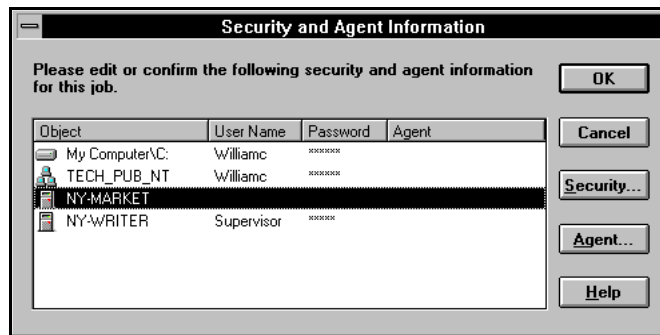


3. Enter a valid username and password.
4. Click OK.

Backing up

You may also change or verify the username and password for a machine or share each time you submit a job. If you have already provided account information for a machine or share, either in the Browse mode or in a previous backup job, ARCserve will retain this information and simply ask you to verify it. (This information is stored for future sessions.)

When you submit a backup job, the Security and Agent Information dialog box opens:



This dialog box displays the security and agent information for each machine or share that is accessed during this job. To change the account information, select the machine/share, then click the Security button.

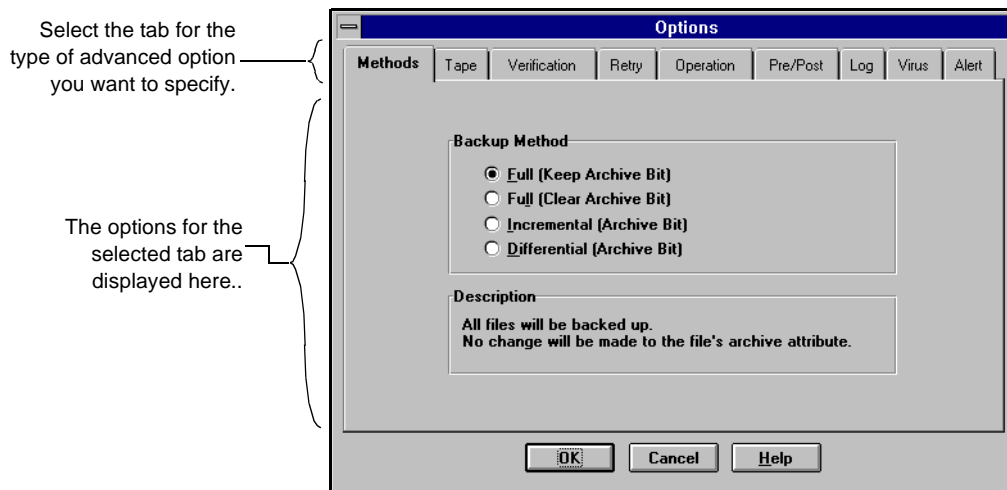


Applying advanced, global backup options



Using the Options dialog box, you can specify options that apply to the entire backup job. To apply options to your job, click the Options button from the Backup Manager toolbar.

The Options dialog box appears:



These global options tell ARCserve how to handle such functions as:

- Backup methods - specify how ARCserve backs up files and whether the archive bit is reset after the file is backed up. For details, refer to the 'Backup methods' section on page 4-18.
- Tape options - specify the tape(s) ARCserve can use and/or overwrite. For details, refer to the 'Tape Options' section on page 4-23.

- Verifying the backup - select how ARCserve will ensure that the data just backed up is readable. Refer to the 'Verifying your backup' section on page 4-28 for details.
- Accessing open files - select how ARCserve backs up files that are already open. For more information, refer to the 'Retry options' section on page 4-30.
- Various utility options - specify whether to record the source file paths in the database, eject the tape from the tape drive, or delete the source files after they are backed up. The 'General backup options' section on page 4-34 discusses these options.
- Pre/post job options - run commands or batch files before the backup job runs and/or after it finishes. Refer to the 'Pre/Post options' section on page 4-37.
- Log options - specify whether ARCserve should keep track of all activity for the job, summary information and errors that occur during the job, or nothing at all. Refer to the 'Log options' section on page 4-38.
- Virus options - enable and configure virus scanning. Refer to the 'Virus scanning options' section on page 4-39.
- Alert options - specify the backup events for which you will be notified and the notification method. Refer to the 'Alert options' section on page 4-40.

The rest of this section describes each backup option in detail.

Backup methods

Backup methods determine what files are backed up and, after the file is backed up, what ARCserve does with the file's archive bit.

What is an archive bit?

Every file has an archive bit. A file's archive bit is a marker that indicates whether the file has been modified. An archive bit that is set (turned on) indicates that the file has changed since it was last backed up or copied.

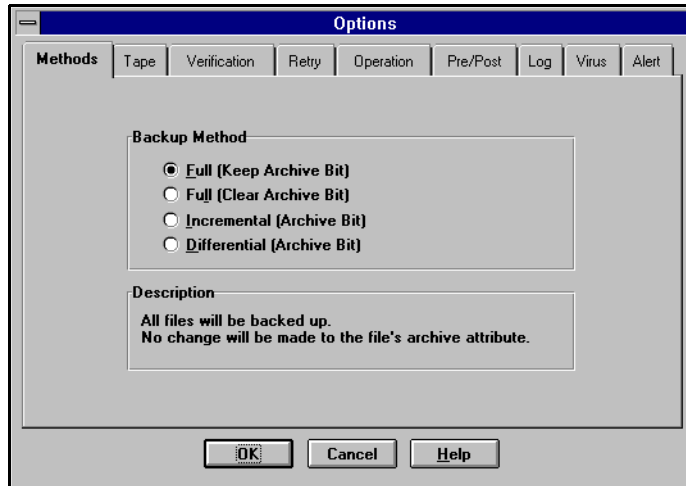
Why turn off a file's archive bit?

Turning off (resetting) the file's archive bit is a procedure that affects what ARCserve will back up the next time. An archive bit that is off indicates to ARCserve that this file has not changed since it was last backed up.

The backup method you select determines whether ARCserve will back up all files, or only files whose archive bit is on. It also determines whether ARCserve will turn off the archive bit after the file has been backed up.



To get to the Methods tab in the Options dialog box, click the Options button on the Backup Manager toolbar.



The following table briefly describes each backup method:

Backup method	Rely on archive bit?	What gets backed up?	Archive bit turned off after the job?
Full (Keep Archive Bit)	No.	All source files.	No.
Full (Clear Archive Bit)	No.	All source files.	Yes.
Incremental (Archive Bit)	Yes.	Source files whose archive bit is turned on.	Yes.
Differential (Archive Bit)	Yes.	Source files whose archive bit is turned on.	No.

Deciding which backup method is right for you

Some of the questions you should be asking yourself when deciding upon a backup method are:

- How much data gets backed up?
- What percentage of my data changes each day?
- How long do I want the backup job to take?

Depending on your needs, each backup method has its benefits. In general, if you have large amounts of data to be backed up every day, running regular differential or incremental backup jobs and an occasional full backup job may be your best solution.



Aside from ARCserve's backup methods turning off a file's archive bit, there are other ways a file's archive bit can be turned on and off:

- Setting/Resetting the file's Archive attribute through another application.
- Running ARCserve copy jobs with a method option that resets the archive bit.

Be careful when selecting a backup method that relies on the archive bit. There may be times when a file is not backed up because another application has changed its archive bit setting.

Using the Full (Clear Archive Bit) method

<hr/> How it works	Full (Clear Archive Bit) backups include all the source files. This method ignores the file's archive bit until after the file is backed up. At that time, ARCserve turns off the file's archive bit. At the end of the job, all files that have been backed up have their archive bits turned off.
<hr/> Benefits	Performing a full backup (using the <i>Full (Clear Archive Bit)</i> method) guarantees that all files selected as your source are backed up at once. This makes it easier for you to restore your data in a crunch and ensures that each backup is up-to-date.
<hr/> Considerations	Full backup jobs are generally time-consuming. If you have a large amount of data to back up, it may be unreasonable to frequently perform backup jobs with the <i>Full (Clear Archive Bit)</i> method option selected.



Using the Incremental method

Incremental backups only include files that have changed since the last *Full (Clear Archive Bit)* or *Incremental* backup (or copy).

<hr/> How it works	These types of backups rely on the file's archive bit. Before backing up a file, ARCserve checks the file's archive bit. If the file's archive bit is on, ARCserve proceeds to back the file up. After the file is backed up, ARCserve turns off the file's archive bit. The next time an incremental backup (or copy) is done, this file is skipped (unless it is modified again).
<hr/> Benefits	In general, these jobs take less time to perform since only files modified since the last full or incremental backup (or copy) are included. This is the quickest method you can use.

<hr/> Considerations	This method requires the most management since only files modified since the last full or incremental backup (or copy) are included. To completely restore your network to the way it was before a disaster, you will need to restore the last full backup and <i>all</i> incremental backups that were performed after that last full backup.
Using the Differential method	
Differential backup jobs include files that have been changed since the last <i>Full (Clear Archive Bit)</i> or <i>Incremental</i> backup (or copy).	
<hr/> How it works	This method relies on the file's archive bit. Before backing up a file, ARCserve checks the file's archive bit. Files whose archive bit is on are backed up. Unlike an incremental backup, however, a differential backup does not turn off the file's archive bit after it has been backed up.
<hr/> Benefits	In general, these jobs take less time to perform than a full backup because only files modified since the last full or incremental backup (or copy) are included. It is also an easier method to manage, since the last differential backup contains all the files that have changed since the last full backup. This means when you want to completely restore your network to the way it was before a disaster, all you need are the last full backup and the last differential backup.
<hr/> Considerations	In general, these jobs take longer to perform than incremental backups. This is because differential backups include <i>all</i> files that have changed since the last full backup.

Tape Options

Tape options apply to the tapes that will be used for the backup job. This section describes each option and how to choose the right one for each job.



To get to the Tape tab in the Options dialog box, click the Options button on the Backup Manager toolbar.

For greater security, provide a password for this backup job.

Specify what tape ARCserve can use to start the job.

Specify what tape(s) ARCserve will use if the job requires more than one tape.

4

Session Password

This option provides greater security for the data you are backing up in this job. The data from this job cannot be restored unless this password is provided.



Use this option with care! If you submit a backup job with a session password, you **MUST** remember the password! None of the job sessions will be restored if you do not provide the correct password in the Restore Manager.

First Tape options

The “first tape” is the tape you use when the job begins. Select the tape you will allow ARCserve to use as the first tape for the backup job.

Append to Tape -
the default
selection

The job session(s) are simply appended to the tape you selected.



Unlike the *Append to Tape* option described above, the following tape options (the “Overwrite...” tape options) format the tape and then start writing data at the beginning of the tape. As a result, any data that had been on the tape is lost.

Also, ARCserve will not overwrite a foreign, non-Cheyenne tape (a tape that has been used by another application). The only way to use a non-Cheyenne tape is to erase the tape and/or format it. Refer to Chapter 9, “Managing Tapes & Tape Devices”, for details about formatting and erasing tapes.

Overwrite Same
Tape Name, or
Blank Tape

Select this option if you want ARCserve to overwrite the tape in the drive *only* if it is the one you specified for the job or if the tape is blank.

Here’s how it works: If the tape in the drive is not the one specified for the job, ARCserve looks to see if it is blank. If the tape isn’t blank either, ARCserve prompts you to supply the

specific tape and waits the time specified in the *Tape Timeout* field for you to do it. When the correct tape is provided, it is reformatted, and the backup job begins.



The job fails if the tape timeout period expires and ARCserve does not find the correct tape.

Overwrite Same
Tape Name or
Blank Tape First,
then Any Tape

Use this option if you want ARCserve to overwrite *any* tape found in the drive.

If you select this tape option, ARCserve checks to see if the tape in the drive is the one specified for the job. If it isn't, ARCserve checks to see if the tape is blank. If the tape isn't blank either, ARCserve simply reformats whatever tape it finds in the device and starts backing up files at the beginning of the tape.

Tape Spanning options

If the backup job you submit requires more than one tape, ARCserve needs you to provide some information on what to do. Whether you have only one or several drives in the tape group, you need to specify which tapes ARCserve can use when the job spans tapes.

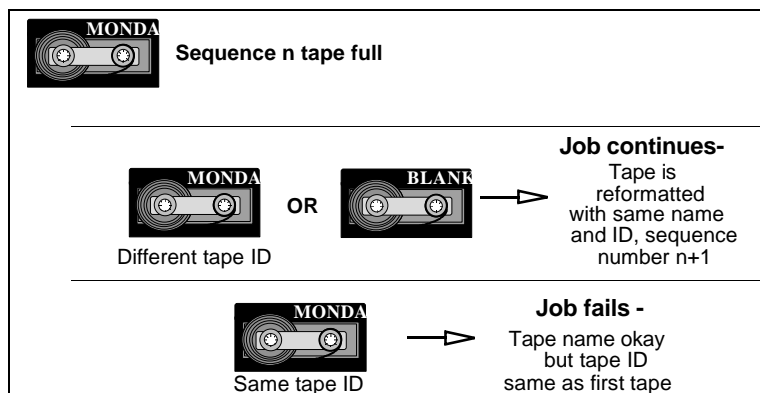


ARCserve will not overwrite a foreign, non-Cheyenne tape (a tape that has been used by another application). The only way to use a non-Cheyenne tape is to first erase and/or format it. Refer to Chapter 9, “Managing Tapes & Tape Devices” for details about formatting and erasing tapes.

Overwrite Same Tape Name, or Blank Tape

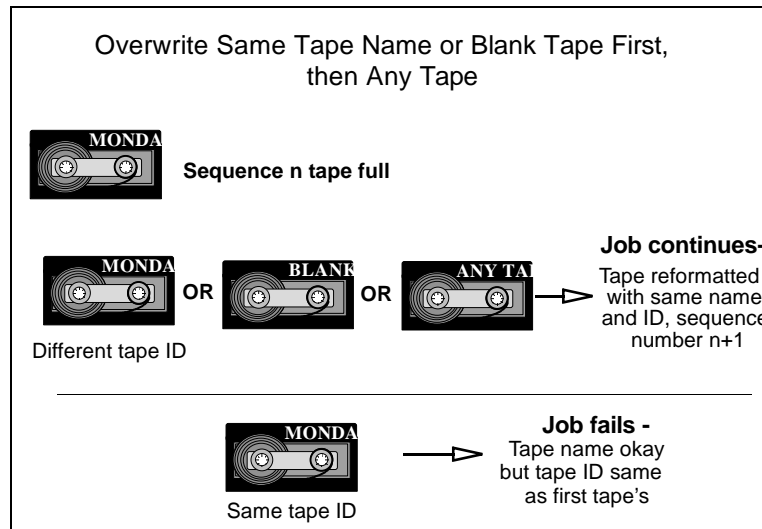
Use this option if you want ARCserve to write to the tape in the device only if it has the same tape name (but a different tape ID) or if it is blank.

ARCserve remembers the name and ID of the job’s first tape. When the job requires an additional tape, ARCserve checks if the new tape has the same name (but different tape ID) or if it is a blank tape. As long as the ID is different, ARCserve reformats the tape, giving it the same name and ID as the first tape. The sequence number will be different.




**Overwrite Same
Tape Name or
Blank Tape First,
then Any Tape**

Use this option to allow ARCserve to overwrite *any* tape found in the device (as long as it has a different ID from the first tape's ID). All subsequent tapes are reformatted with the same name and ID as the first tape. Only the sequence number will be different.



Verifying your backup

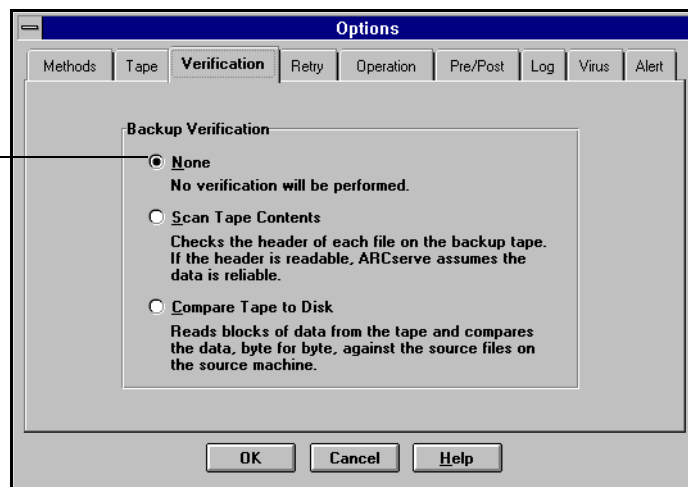


As an added measure of ensuring data integrity, ARCserve provides you with two methods of verifying that the data you just backed up to the tape is reliable. Since these options require performing another pass over the tape (which may be time consuming), ARCserve, by default, does not perform either of these options. However, it is a good idea to select one of these options to ensure the most dependable backups.



To get to the Verification tab, click the Options button on the Backup Manager toolbar.

By default, ARCserve does not perform a verification function.



The job verification methods are:

- Scan Tape Contents
- Compare Tape to Disk


Scan Tape Contents

If you select *Scan Tape Contents*, ARCserve checks the header of each file on the backup tape. If the header is readable, ARCserve assumes the data is reliable. If the header is not readable, the Activity Log is updated with this information. Of the two methods, this one is faster.

Compare Tape to Disk

If you select *Compare Tape to Disk*, ARCserve reads blocks of data from the tape and compares the data, byte for byte, against the source files on the source machine. This option is more time consuming, but it ensures that all data on the tape is *exactly* as it is on disk. If ARCserve finds a mismatch, the Activity Log is updated with this information.

Retry options



ARCserve needs to know how to handle files when backing them up. It needs to know two things:

- Open File Retry options - how to back up files that are already open by another process.
- File Sharing options - if ARCserve gets to the file before another process, can other processes access it while it is being backed up.



To get to the Retry tab, click the Options button on the Backup Manager toolbar.

The screenshot shows the 'Options' dialog box with the 'Retry' tab selected. The dialog has several tabs: Methods, Tape, Verification, **Retry**, Operation, Pre/Post, Log, Virus, and Alert. The 'Open File Retry' section contains the following options:

- Retry Immediately
- Retry After Job
- Number of Retries: 3
- Retry Interval (sec): 10

The 'File Sharing' section contains the following options:

- Use Deny None if Deny Write Fails
- Use Lock Mode if Deny Write Fails
- Deny Write
- Deny None


At the bottom are buttons for OK, Cancel, and Help. Two text annotations with arrows point to the 'Open File Retry' and 'File Sharing' sections:

- When and how often ARCserve will try again to back up a file that was open when it first tried to back it up.
- Whether and how other processes can access the file as ARCserve is backing it up.

Open File Retry options

If, during the backup, any files were unavailable to ARCserve because they were locked by another process (already opened and being used by another application), ARCserve can attempt to back them up again by:

- Trying again immediately, or
- Trying again after the job



You don't have to limit yourself to one selection here. This option allows you to select none, one, or both methods.

Retry Immediately

If you select *Retry Immediately*, ARCserve will try to back up 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

If you select *Retry After Job*, ARCserve will try to back up 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


Specify the number of times you want ARCserve to try to back up the file.

Retry Interval

Specify a period of time you want ARCserve to wait between attempts.

File Sharing option

When ARCserve attempts to open a file to back it up, it needs to know whether you want subsequent users and processes to be able to access this file and, if they can access it, whether they can save data to it. If the file is already open by another process and ARCserve cannot open the file with the selection you make here, ARCserve will then use the *Open File Retry* option (discussed on the previous page) that you specified to attempt to back up the file again.



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 discussed below. As long as no other process is writing to these files during the job, the backup will be consistent.

If you want to ensure that only the most current version of every file is backed up, you should select a *Deny Write* or *Lock Mode* option discussed below.

Use Deny None if Deny Write Fails
- the default selection

This is the default option ARCserve uses when accessing files. ARCserve attempts to place the file in “Deny Write” mode (which allows users to open and view, but not write to, the file). If ARCserve can’t do this (because the file is already open), ARCserve places the file into “Deny None” mode (which lets users open and write to the file), and then backs it up.

*Use Lock Mode if
Deny Write Fails*

ARCserve attempts to place the file in “Deny Write” mode (which allows users to open and view, but not write to, the file). If ARCserve can’t do this (because the file is already open), it locks the file completely (prohibiting any user from opening or writing to the file) and backs it up. This method ensures the most recent version of the file is backed up.

Deny Write

Deny Write means that while ARCserve has the file open and is backing it up, another process can only open the file for reading.

A file can only be opened in *Deny Write* mode if it is not already opened by another process. This mode is the safest for backups, as no other process can write to the file during the backup job. If, however, another process has the file open before ARCserve gets a chance to open it, ARCserve will not be able to open the file, and the file will not be backed up (unless you specified an *Open File Retry* option, which was discussed in the previous section).

Deny None

Deny None mode will allow another process to read or write to the file, regardless of whether ARCserve opens it first or opens it after another process already has it open. This method ensures that your files are up-to-date, although the file that was backed up to tape may not be.

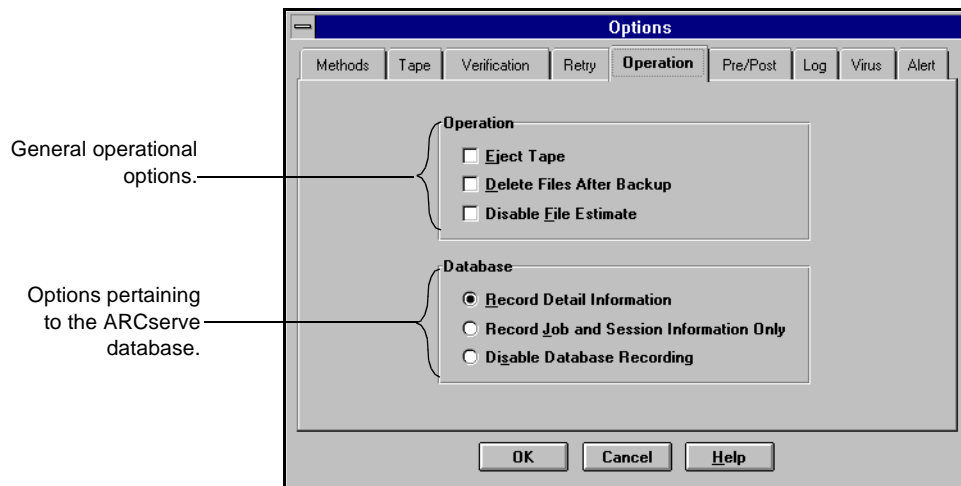
General backup options

There are several general options you can specify for your backup jobs. These include:

- Not recording any information about the job into the ARCserve database
- Ejecting the tape after the job is completed - this option functions properly only on tape devices that support this feature
- Deleting files from the source disk after they have been backed up




To get to the Operation tab in the Options dialog box, click the Options button on the Backup Manager toolbar.



Each option is discussed on the following pages.

Operation options



Eject Tape - ensure the tape is not accidentally used for another job.

Set this option to have ARCserve eject the tape from the drive after the job finishes. This helps prevent any other job from overwriting information on this tape.

Delete Files after Backup - use this option with filters to groom your disks

Select this option to delete source files from the source machine after they have been backed up to tape. Use this option as a way of performing disk grooming. Set a filter to back up files that haven't been accessed for a certain period of time (say, 3 months or so), and then select this option to delete these old, unaccessed files from the source disk. Refer to Chapter 7, "Customizing Your Jobs", for more information about filters.

The source files are deleted when the session is completed. If you select the "Log All Activity" Log option, all the deleted files will be listed in the Job Log. Refer to the 'Log options' section on page 4-38 for more information on the "Log All Activity" option.

Disable File Estimate

Before any file is backed up to tape, ARCserve, by default, performs an estimate of how long the job will take. Set this option on if you want ARCserve to skip this function.



Database options



Record Detail Information

By default, this option is selected. All the information pertaining to the source and destination are logged in the database.

If you do not want details about this job recorded in the ARCserve database, select one of the Database options discussed below. You may want to do this when you have limited disk space.

Record Job and Session Information Only

Select this option if you want only job and session information logged in the ARCserve database.

Disable Database Recording

Select this option if you don't want any information about this job logged in the database. Select this option if you are backing up the ARCserve database or if you have limited disk space.



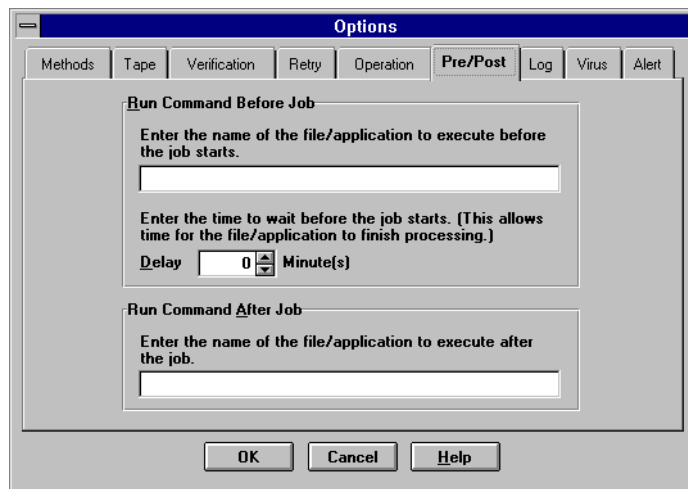
If you select this option, you will have to submit a Merge Tape operation before you can use any of the database views to restore this data. Or, you can restore data from this job using the Restore by Tape Media source view.

Pre/Post options

These options allow 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 backed up and then run a batch file you created that sends a detailed report to the printer.



To get to the Pre/Post tab in the Options dialog box, click the Options button on the Backup Manager toolbar.



Run Command Before Job

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

Delay n Minutes

Enter the time to wait (in minutes), before the backup job will start. This period of time gives the application or file specified in the *Run Command Before Job* field time to process.

Run Command After Job

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

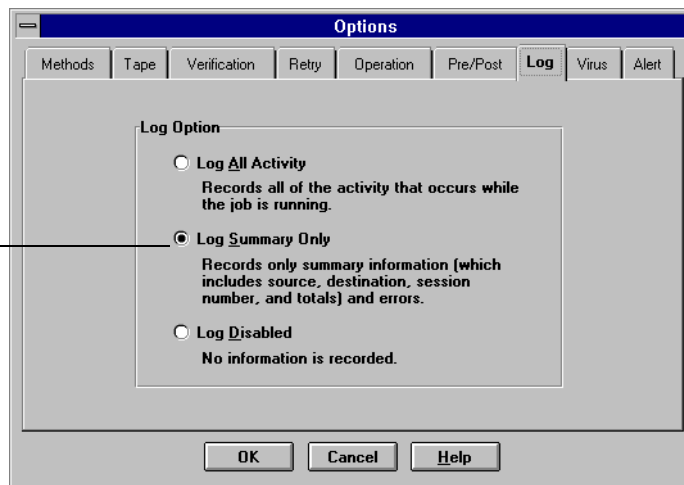
Log options

Select the Log tab in the Options dialog box to determine the level of detail you will see in the Job Queue Log or Job View Database Log.



To get to the Log tab in the Options dialog box, click the Options button on the Backup Manager toolbar.

By default, a summary of the job is recorded in the Job Log.



Log All Activity

If you select this option, all of the activity that occurred while the job is running will appear in the Job Log.

Log Summary Only (the default selection)

If you select this option, you will get a summary of the job (which includes source, destination, session number, and totals) along with any errors that may have occurred.

Log Disabled

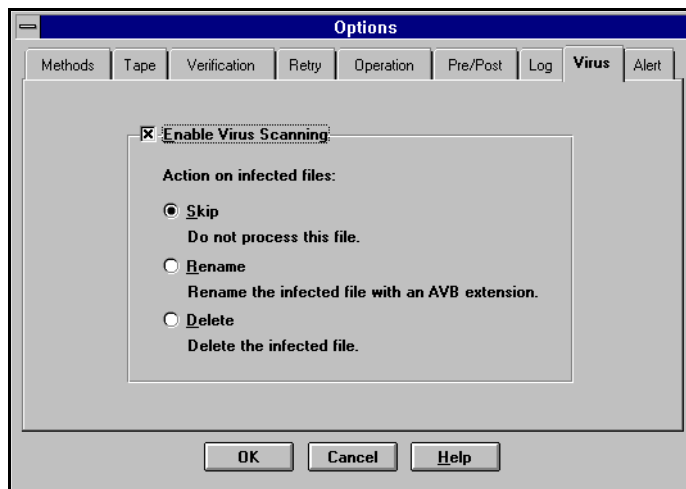
Select this option when you do not want any messages about this job saved to the Job Log.

Virus scanning options

You can automatically scan for viruses during a backup operation using the virus scanning options.



To get to the Virus tab in the Options dialog box, click the Options button on the Backup Manager toolbar.



Enable Virus Scanning

Select this option to enable virus scanning.

The table below describes the virus scanning options:

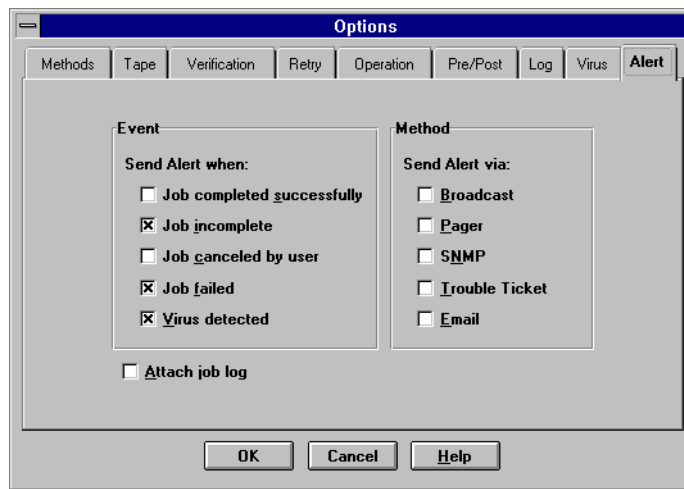
Option	Description
Skip	Do not back up the infected file.
Rename	Rename the infected files with the extension AVB. (If a file with the same name and the extension AVB exists, then extension AV1 is used, and so on.)
Delete	Delete the infected file.

Alert options

You can use the Alert notification system to send messages about events in your backup operation. For more information on the Alert system, refer to Chapter 16, “Using the Alert Manager”.



To get to the Alert tab in the Options dialog box, click the Options button on the Backup Manager toolbar.



Event

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

Event	Description
Job completed successfully	All of the nodes and drives/shares were processed.
Job incomplete	Some nodes and/or drives/shares were missed.
Job cancelled by user	The user cancelled the job.
Job failed	The job was started but could not be completed.

Event	Description
Virus detected	A virus was detected in one of the files to be backed up. Refer to the 'Virus scanning options' section on page 4-39 for more information.

Method

Choose one or more of the following Alert notification methods. Before you choose a method, make sure that the Alert system is properly configured for this method. Refer to Chapter 16, "Using the Alert Manager" for more information.

Method	Description
Broadcast	Windows NT Broadcast's can be sent to specific users or groups.
Pager	Alphanumeric paging system.
SNMP	Simple Network Management Protocol (such as NetWare Management System (NMS) and HP Open View).
Trouble Ticket	Any print queue on your network. (Trouble Tickets can only be sent to non-Postscript printers.)
Email	Microsoft mail.



Attach job log

Select this option to include the job log information in the Alert message. (This option applies for Email only.)



5

Chapter

RESTORING YOUR DATA

This chapter contains information about restoring data to local Windows NT machines and shared network drives.


In this chapter, you will learn:

<i>Page</i>	
5-4 ➤	How to perform a quick, basic restore
5-8 ➤	How to change and use the different restore source views
5-24 ➤	How to use the advanced restore options

Restore basics

Source comes from a tape, destination is a hard drive

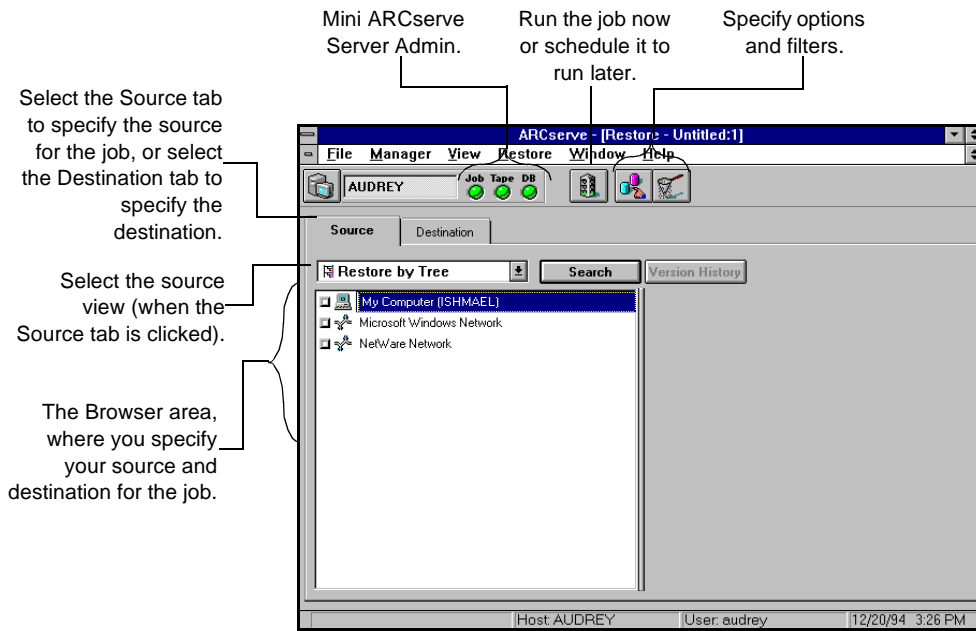
Restoring data back to your machines on the network is easy and flexible with ARCserve. Each restore job requires a source and destination. The files you select as your source will always come from a tape, and the place you select as your destination will always be a hard drive.



ARCserve provides you with four methods for selecting the data you want to restore (the source). Three of these methods use information logged in ARCserve's database, while the fourth method doesn't. You choose the method that works best for you.

You also have several methods for restoring the data to the machine. Start restoring data at a drive level or dive deeper into the machine's directory tree and specify a subdirectory from which ARCserve should start restoring to. You can choose to recreate the entire path that was backed up, or only from the base directory level.

The Restore Manager is shown below:



Restoring files - the basic steps

The following steps outline a basic restore job - that is, using the default restore options to get back data that has been backed up to tape.



The following basic steps assume that your ARCserve database is up-to-date and available. Although the following example uses the “Restore by Tree” source view, the basic steps can be applied using any of the source views. Refer to the ‘Changing the source view’ section on page 5-8 for more information on using a different source view.

Someone accidentally deleted an important directory from Sean’s machine. Since this directory was recently backed up, he will restore the directory and its contents back to his hard drive.

Steps for a basic restore job

Follow the steps below and apply the example above to learn the basic steps for restoring data.



1. Click the Restore button from the Quick Access window .
The Restore Manager appears, with the default source view, Restore by Tree, selected.

Refer to the 'Changing the source view' section on page 5-8

2. Select what you want to restore.
Select the Source tab and choose the view you want to use to select the source for your job. Then specify what you want to restore.

Make sure the Source tab is selected...

...and choose the source view you want to use. In this case, it is Restore by Tree.


After finding what you want to restore, click the green box next to it to select it.

The screenshot shows the ARCserve Restore window with the 'Source' tab selected. The 'Restore by Tree' view is active, displaying a file tree under 'My Computer (ISHMAEL)'. The 'MARKET' folder is selected, and a green checkmark is visible next to it. The right-hand panel displays tape metadata: Tape Name (PROJECTS_NOV1994), Tape ID (7558), Sequence No. (1), Session No. (5), Backed Up On (11/08/94 13:49:14), and Source Path (D\$).

3. Specify where you want the files restored to.
Select the Destination tab.

Select the Destination tab...

The screenshot shows the ARCserve Restore window with the 'Destination' tab selected. The checkbox 'Restore files to their original location(s)' is checked. The file tree shows 'My Computer (TECH_PUB_SERVER)' selected. The right-hand panel displays system information: CPU Type (486), Number of CPU's (1), Total Memory (16,711,680), and Free Memory (802,816).

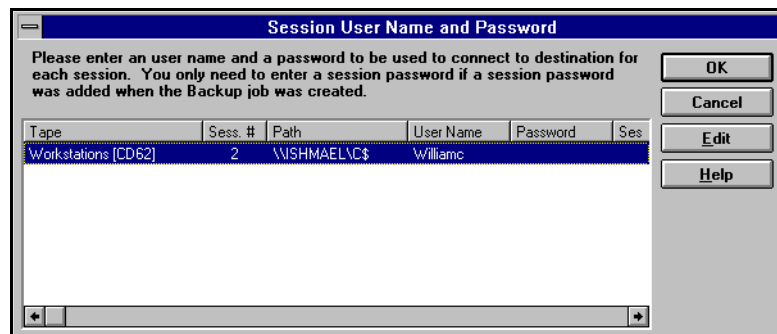


The option *Restore files to their original location(s)* is selected by default. When this option is selected, your files are restored to their original machine and path. If the original location is on a remote machine, you will be prompted to enter a valid username and password for the remote machine.

You can also clear this option and select another location using the Browser.

4. Click the Run/Schedule button.

The Session User Name and Password dialog box opens:



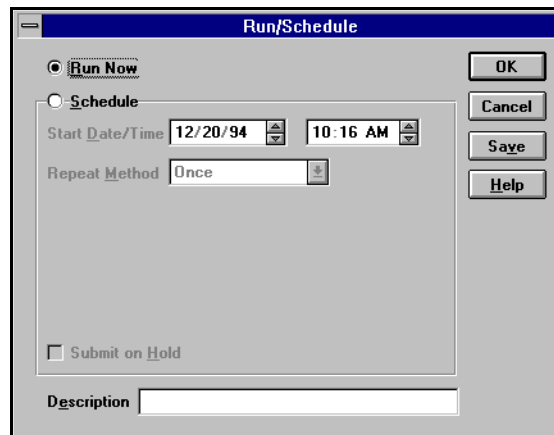
Enter, change, or verify the username and password for the destination machine for each tape session. If a session password has been assigned, enter the session password as well. For more information on session passwords, refer to Chapter 4, “Backing Up to Tape”.

To enter or change the security information for a session, select the session and click the Edit button.

Refer to Chapter
7, “Customizing
Your Jobs”

5. Click OK to continue.

The Run/Schedule dialog box opens:



The Run/Schedule dialog box enables you to run the job immediately, or schedule the job for a later time.

To save this job to a script file, click the Save button in this dialog box.

6. Click OK to submit this job.

Changing the source view

Although the files that you restore will always come from a tape, the way you select these source files can vary. ARCserve provides you with the most convenient and thorough methods for finding the files you want to restore. There are four methods by which you can choose your source files:

- The Restore by Tree view - Using its database, ARCserve reconstructs a directory tree in the Browser of all machines, directories, and files that were backed up.
- The Restore by Tape view - Using its database, ARCserve displays all tapes used for backups and the source for each session on the tape.
- The Restore by Query view - You specify a search pattern and ARCserve searches its database for the specified files and directories.
- The Restore by Tape Media restore - Place a tape in the tape device and select a particular tape session to restore.

The method you choose to select your source depends on what you know about the files you want to restore and the tape(s) you need to use.

The ARCserve database provides a view of what exists on all tapes

The four available restore methods display the source information either by using the ARCserve database or by reading what tape devices are currently attached to the ARCserve host. “Restore by Tree”, “Restore by Tape”, and “Restore by Query” use the database, “Restore by Tape Media” does not.

When using “Restore by Tree”, “Restore by Tape”, or “Restore by Query” source views, you are using the ARCserve database. The information you use to select your source is based on all

the backups made since you started using ARCserve. Unless you have pruned your database, you should be able to restore all files from all backups you have performed with ARCserve.

“Restore by Tree” displays a graphical tree of machines that have been backed up

“Restore by Tree” is the default source view for the Restore Manager. All the machines that you have backed up appear in the Browser. Selecting a machine displays the drives, directories, and files that were backed up from that machine. If you have backed up something several times (the ARCserve home directory on your local machine, for example), it is displayed only once in the graphical tree. By default, the latest version of a file is restored, but you can restore an older version by using the Version History feature.

Use “Restore by Tree” when you don’t know which tape contains the data you need, but you know which machine it came from. This source view provides you with the best overall picture of your network and each machine as it was when you backed it up.

“Restore by Tape” displays a graphical tree of backup tapes and sessions

The “Restore by Tape” source view displays a list of all the tapes that ARCserve has used to back up data. Selecting a tape displays a list of sessions on that tape and what was backed up for that session. All information displayed is gathered from ARCserve’s database.

This source view is a good one to choose if you know which tape(s) you used for the backup, but you’re not sure which session(s) contains the data you need to restore.

“Restore by Query” displays files or directories based on a search pattern that you specify

The “Restore by Query” source view uses the ARCserve database and a search pattern that you specify to display files you want to restore. The Restore by Query can be as general or specific as you want it. You can specify a search pattern directed towards a specific machine, directories, or files. Subdirectories can also be included in the search.

Use the “Database Restore by Query” source view when you know the names of the files and/or directories you want to restore, but do not know the machine it was backed up from or the tape it was backed up to.

“Restore by Tape Media” displays tape groups and devices attached to the host

The “Restore by Tape Media” restore method displays all the device groups attached to your machine and the tapes in them. You select the device group and tape, and then you specify which session(s) on the tape you want to restore.

Since this source view does not use the ARCserve database, you cannot pick and choose files and directories you want to restore. Unless file or directory filters are applied to the job, every file in the session you specify as the source is restored.

If you know the exact tape and session number to use, this source view is an effective and easy way of performing a quick restore. This source view is also useful for disaster recovery, when the ARCserve database is not available or the tape’s history has not been merged into the ARCserve database.

Using the Restore by Tree source view

“Restore by Tree” is the default source view for the ARCserve Restore Manager. This view displays a graphical tree of files and directories on Windows NT machines and NetWare servers that were backed up with ARCserve. After you specify the files and directories you want to restore, ARCserve lists the tapes you need for the job.



This source view is useful if:

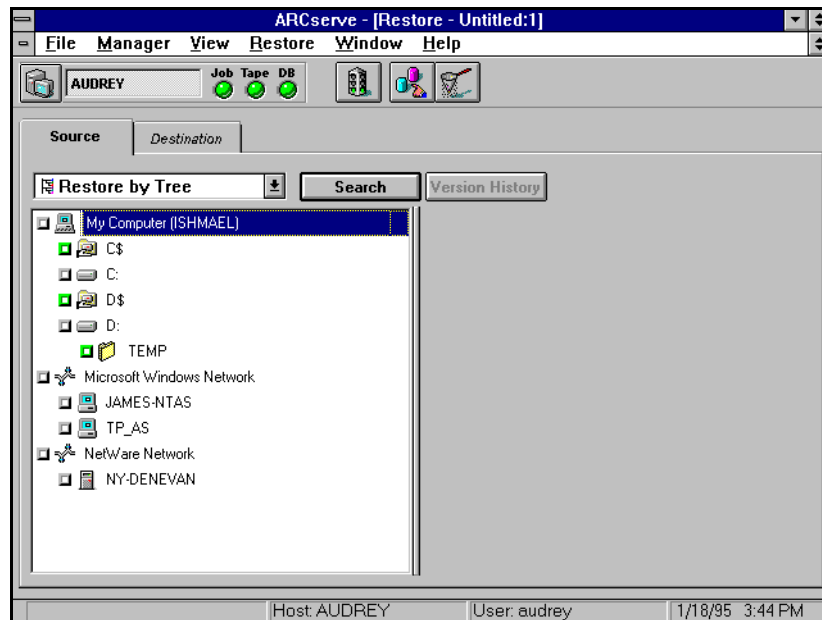
- You want to restore a directory or drive that was backed up from a particular machine back to that machine.
 - You want to restore a directory or drive that was backed up from a particular machine to another machine.
-

Follow these directions to use the “Restore by Tree” view to select what you want to restore:

1. In the Restore Manager, select Restore by Tree from the source view drop-down list.

The Browser displays a graphical tree of the networks and machines that have been backed up by ARCserve. Paths for files and directories that have been backed up appear when you expand the tree for each machine, allowing you

to select your source based upon the machines the data was backed up from.



2. Select the files you want to restore.
Refer to Chapter 3, "Basic Procedures", for details on selecting files and directories from machines.

Version History and Search buttons help you find and select files

The “Restore by Tree” source view provides you with two features that help you locate and select specific files or directories. They are the:

- Version History button - select an older backup of the file or directory to restore
- Search button - specify a search pattern and ARCserve returns all files and directories that match it




You can only use the Version History and Search features when you are using “Restore by Tree” as your source view.

Version History feature

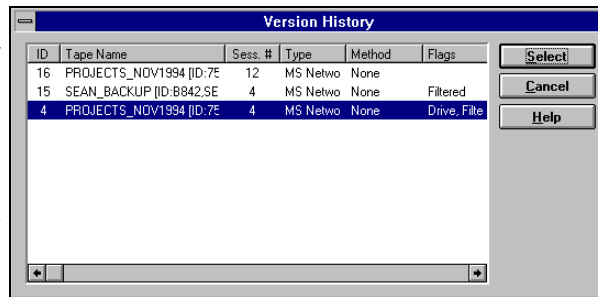
Use the Version History feature to help you select the version of the file, directory, or drive you want to restore.

Sean wants to restore an older version of his directory, but he's not sure which tape contains a backup of this directory.

1. In the left pane of the Source area, select the drive, directory, or file you want to restore.
2. Click the Version History button. 

ARCserve searches its database and returns a list of all the backed up versions of this file, directory, or drive. From this list, you can select the version you want to restore:

A list of all versions of the volume, drive, directory, or file appears. By default, the latest version will be selected. Highlight the version you want, then click the Select button.



If you try to expand a backup session that has been pruned from the database, you will be asked whether you want to merge the pruned session. Refer to Chapter 14, "Adminstrating the ARCserve Server" for more information on pruning. Refer to Chapter 11, "Using the Utilities" for more information on merging.

Search feature

Use the Search feature to have ARCserve quickly and automatically locate files and directories whose name or location you are unsure or unaware of.


Sean needs to recover an old report document, but he can't remember the name of the report or on whose machine it was saved to. The only information he knows about the file is that it has a .DOC extension.

1. From the Restore Manager, click the Search button. 

The Search dialog box appears.

2. Enter the search criteria and click Search.

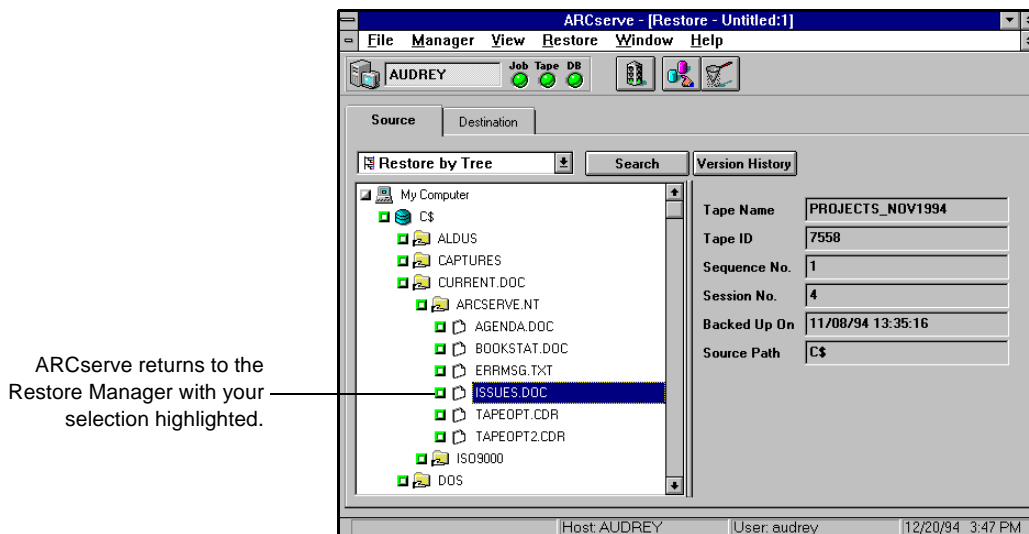
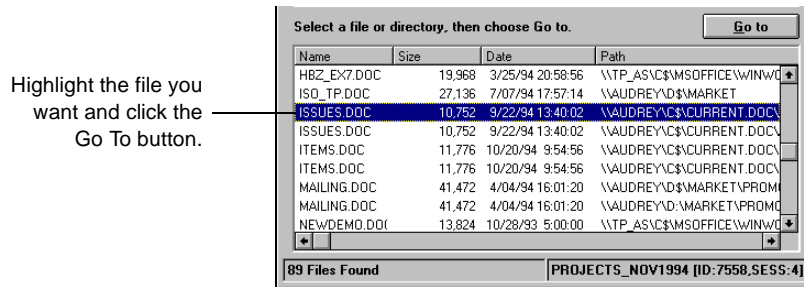
Use an asterisk (*) or question mark (?) as a wildcard symbol, as necessary.



After you click the Search button, ARCserve searches its database, and displays the files and directories that match the search pattern you specified.

3. Highlight the file you want and click the Go To button .

The Restore Manager reappears, as ARCserve expands the directory tree and highlights the file. At this point, you can select it for restoring.



Using the Restore by Tape source view

The Restore by Tape source view allows you to select files to restore from all the tapes that have been used for backup jobs. With this view, you first select the tape you need for the restore job, then the session, and finally, the files and directories you want to restore. You can also use this method to restore entire sessions on the tape.

This source view uses ARCserve's database. As long as the information about the tapes you have used for backups is in the database, you can use this method effectively for selecting files to restore.



This source view is useful if:

- You know the tape that holds the data you want to restore, but you may not be sure which session it is in.
-

Sean maintains backups on tapes based on the months of the year. He wants to restore files from a project he worked on in November, but he can't remember which session these files are in.

5

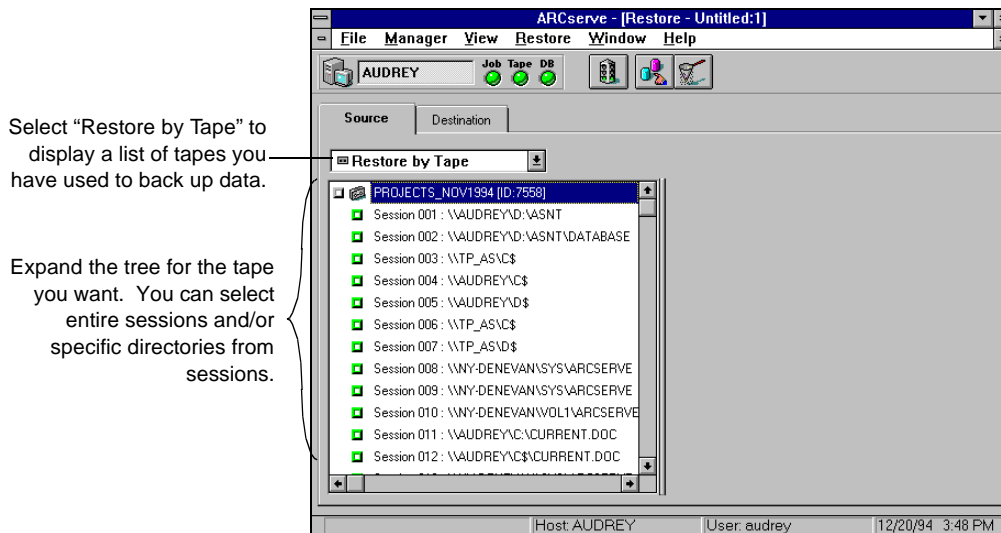
The steps below detail how to use the Restore by Tape view to restore:

1. In the Restore Manager, select Restore by Tape from the source view drop-down list.

A list of tapes that you have used for backups displays in the Browser.

2. Expand the tree for the tape you want to restore from .

Double-click on the name of the tape that holds your data to display a list of sessions logged in the database for that tape.



If you try to expand a backup session that has been pruned from the database, you will be asked whether you want to merge the pruned session. Refer to Chapter 14, "Administrating the ARCserve Server" for more information on pruning. Refer to Chapter 11, "Using the Utilities" for more information on merging.

3. Select or expand the tree for the session that contains the files you need to restore.

If you select the session (click the box next to the session name), ARCserve will restore every file from that session.

If you only want to restore specific files or directories from a session, expand the tree for the session, and then select the files and directories.

4. Repeat step 2 and 3 for each session on the tape you want to restore from.

Using the Restore by Query source view

Use the Restore by Query source view to restore files 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 tape(s) it was backed up to.

This view also uses ARCserve's database. As long as the file and directory information is logged in the database, you can effectively use this view to select files to restore.

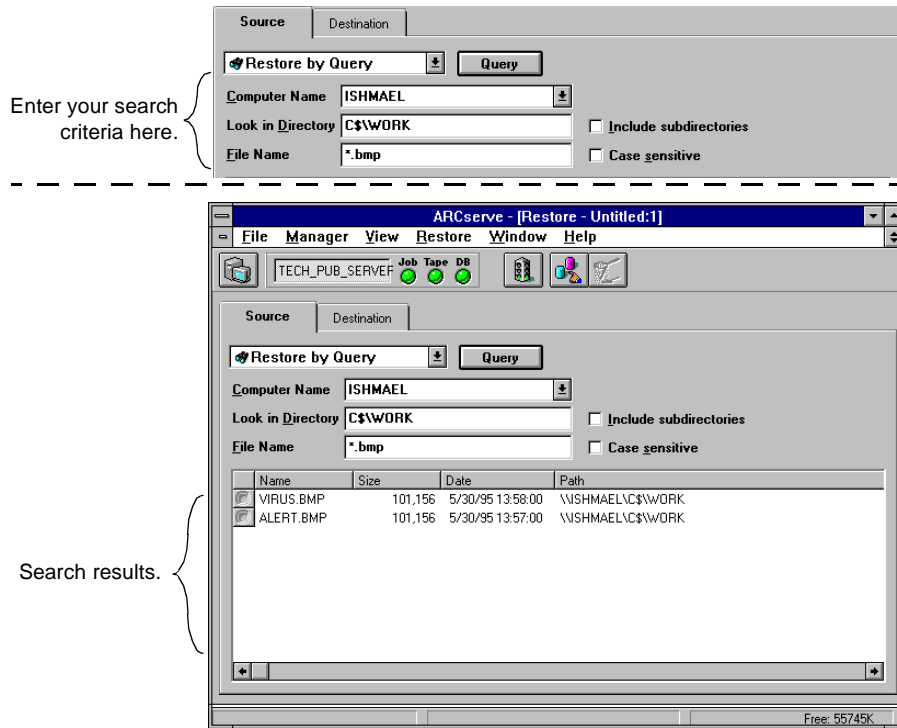


This source view is useful if:

- You know the entire path, a partial path, or only the name of the file or directory to restore.
- You don't know which machine it was backed up from or which tape it was backed up to.

-
1. In the Restore Manager, select Restore by Query from the source view drop-down list.

2. Specify your search criteria and click the Query button .



3. Select the file(s) and directory(ies) you want to restore .
Double-click the name of the file or directory. A green light appears when a file or directory is selected.

Using the Restore by Tape Media source view

Use the Restore by Tape Media source view to restore complete sessions on a tape in a tape device. All files in the session are restored to the destination, unless filters are added to the job.

This view does not rely on the ARCserve database. Instead, a real-time view of the tapes that are in your tape devices is shown. You specify a session number and ARCserve restores all the files from that session.



This source view is useful if:

- Information from your ARCserve database is unavailable (possibly because the Database Engine is stopped, your database is damaged, or the information about the files or directories you need were not included when you backed them up).
 - You are trying to recover your network after a disaster (disaster recovery).
-

1. In the Restore Manager, select Restore by Tape Media from the source view drop-down list.
ARCserve will read directly from the tape devices and display your tape groups and tapes in the Browser.
2. Select the tape group and tape you want to restore from.
Refer to Chapter 3, “Basic Procedures”, for details on using the Browser to select a tape group and tape.
3. Specify the session number to restore.

In the *Sess.* field, specify the number of the session you want to restore.

The screenshot shows the ARCserve Restore dialog box. The 'Source' tab is active, and the 'Restore by Tape Media' option is selected. The 'Group' is set to 'GROUP0', the 'Tape' is 'ACCTG3', and the 'Sess.' field is set to '4'. The 'Device Information' section shows 'Name: TANDBERG TDC 4200' and 'Config: Device:2 = Board:0, Bus:0, SCSI ID:4.'. The 'Tape Information' section shows 'Name: ACCTG3', 'Sequence: 1', 'Write Protect: No', 'ID: F5A0', and 'Expiration Date: 11/03/95 1'. The status bar at the bottom indicates 'Host: TP_AS', 'User: Administrator', and '12/20/94 4:52 PM'.

Select "Restore by Tape Media",,,

...then select the tape group and tape you want to restore from.

Specify the session number you want to restore. In this case, the fourth session on the tape is selected.

Using advanced options

ARCserve provides several types of advanced restore options. They are:

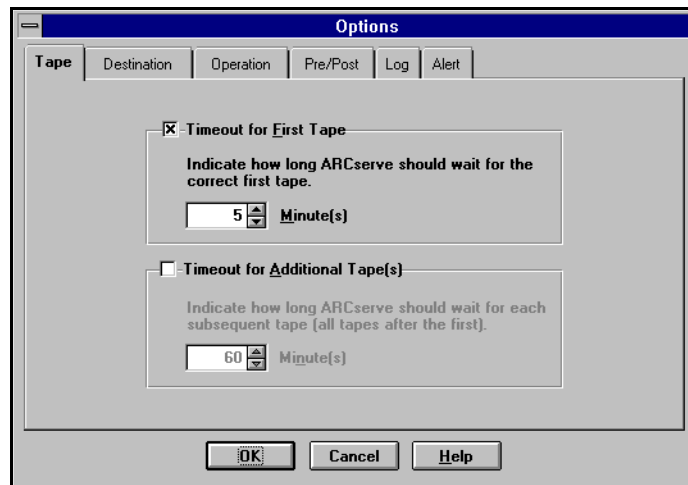
- Tape options
- Destination options
- General operation options
- Pre/Post Execution commands
- Log options
- Alert options

To add restore options to a job:

1. Click the Options button on the toolbar.



The Options screen appears, with the Tape tab selected by default:



2. Select the tab for the type of option you want to apply to the job.

Tape Options

The following options allow you to specify a timeout period ARCserve should wait for you to provide the tape you need to restore your data.



To get to the Tape tab in the Options dialog box, click the Options button on the Restore Manager toolbar.

The screenshot shows the 'Tape' tab of a dialog box. It contains two main sections:

- Timeout for First Tape:** This section has a checked checkbox. Below it is a text box that says 'Indicate how long ARCserve should wait for the correct first tape.' and a spinner control set to the value '5' with the unit 'Minute(s)'.
- Timeout for Additional Tape(s):** This section has an unchecked checkbox. Below it is a text box that says 'Indicate how long ARCserve should wait for each subsequent tape (all tapes after the first).' and a spinner control set to the value '60' with the unit 'Minute(s)'.

Timeout for First Tape

Specify a period of time that ARCserve should wait for the first tape required for your restore job. If the time expires and the first tape has not been provided, the job will fail. By default, ARCserve will wait five minutes.

Timeout for Additional Tape(s)

Sometimes, your data will come from several tapes. By default, ARCserve will not time out. This means ARCserve will continue to wait for you to provide any additional tapes required until you cancel the job. Select this option to specify a timeout period ARCserve should wait. If you select this option, the default timeout period is 60 minutes.

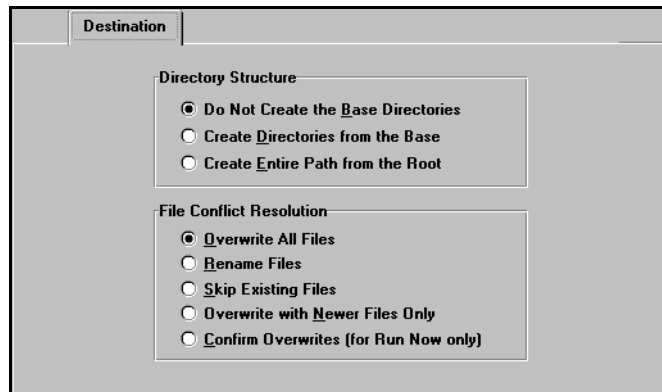
Destination options

ARCserve provides two destination options:

- How to create the directory structure on the destination
- How to handle files that may already exist on the destination



To get to the Destination tab in the Options dialog box, click the Options button on the Restore Manager toolbar.



Destination

Directory Structure

- Do Not Create the Base Directories
- Create Directories from the Base
- Create Entire Path from the Root

File Conflict Resolution


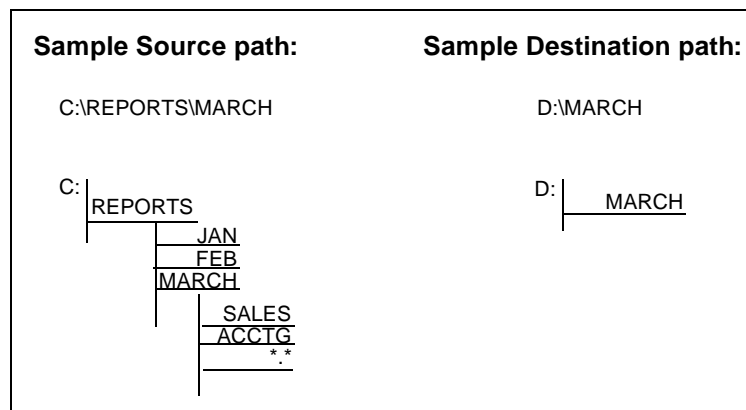
- Overwrite All Files
- Rename Files
- Skip Existing Files
- Overwrite with Newer Files Only
- Confirm Overwrites (for Run Now only)

Directory Structure

Select the method ARCserve should use to create directories on the destination that you select. The default is *Do Not Create the Base Directories*.

Suppose you want to restore the path C:\REPORTS\MARCH. There are several files directly under the MARCH directory, and there are two subdirectories, \SALES and \ACCTG. You want to restore to your MARCH directory on your D drive.

Source path and destination path for Directory Structure options examples




Do Not Create the Base Directories
- the default selection

A “base directory” is considered the last directory selected in the source path. This option will not create the base directory on the destination path, but will create all subdirectories that may reside below the source base directory.

Using this option to restore C:\REPORTS\MARCH to D:\MARCH would produce the following:

Since MARCH is the base directory, it is not created on the destination. However, all files and subdirectories below it are created.




Create Directories from the Base

This option will create the destination path beginning from the base directory.

Using this option to restore C:\REPORTS\MARCH to D:\MARCH would produce the following:

Using the above example, MARCH is the base directory. As a result, everything on the source beginning with the MARCH directory is created on the destination.




Create Entire Path from the Root

This option will create the entire source path (except the root drive or volume name) on the destination. No files from any parent directories are restored. Only the directory path to the base directory is created on the destination.

Using this option to restore C:\REPORTS\MARCH to D:\MARCH would produce the following:

Using the above example, C: is the root. As a result, the path to the MARCH directory is built beginning from the REPORTS directory. Note that no files from parent directories are restored - only the path to the base directory is created.



File Conflict Resolution

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. The default is *Overwrite All Files*.

Overwrite All Files

This option will restore all source files to the destination regardless of conflicting file names. Existing files on the destination will be overwritten by the files from the source.

Rename Files

This option copies 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. Filenames without extensions will be renamed with extensions “.__0”, “.__1”, “.__2”, etc.

Skip Existing Files

If you select this option, ARCserve will not restore a source file if a file with the same name already exists on the destination.

Overwrite with Newer Files Only

If you select this option, ARCserve will only restore source files whose modification date is later than the modification date of the file with the same name on the destination. Source files whose modification date is earlier will not be copied to the destination.

Confirm Overwrites

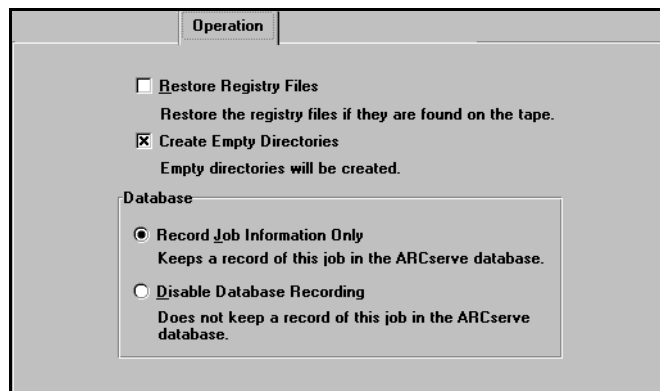
This option is only applicable when submitting an immediate job (a job with Run Now selected in the Run/Schedule dialog box). Before ARCserve tries to restore the source file, you will be prompted to confirm that you want the file on the destination to be overwritten.

General operation options

Following are some general options that can be applied to your job.



To get to the Operation tab in the Options dialog box, click the Options button on the Restore Manager toolbar.



Restore Registry Files

Select this option if you want registry files restored to the machine.

Create Empty Directories

Select this option if you want restore directories that have no files in them.

Database options

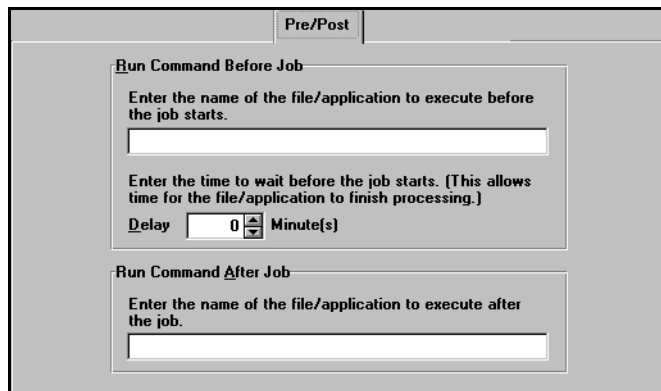
These options decide whether ARCserve will record the restore job in its database. One of these options is always selected for restore jobs. By default, the *Record Job Information Only* option is selected. If you do not want this job recorded in the database, select the *Disable Database Recording* option.

Pre/Post options

These options allow you to run a command on your machine before and/or after the restore job is executed.



To get to the Pre/Post tab in the Options dialog box, click the Options button on the Restore Manager toolbar.



Run Command Before Job

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

Delay n Minutes

Enter the time to wait (in minutes), before the restore job will start. This period of time gives the application or file specified in the *Run Command Before Job* field time to process.

Run Command After Job

Enter the name of the file to be executed on the machine after the job has completed.

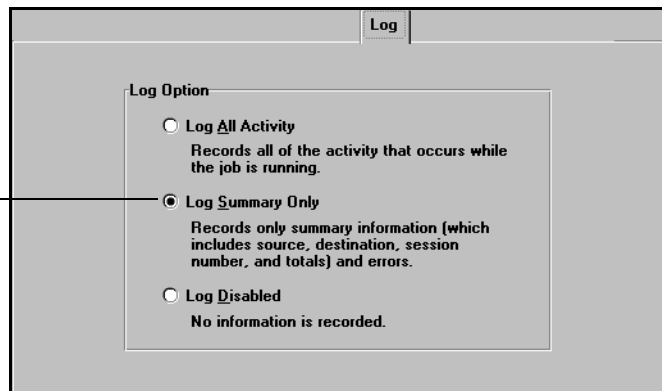
Log options

Select the Log tab in the Options dialog box to determine the level of detail you will see in the Job Queue Log or Job View Database Log.



To get to the Log tab in the Options dialog box, click the Options button on the Restore Manager toolbar.

By default, a summary of the job is recorded in the Job Log.



Log All Activity

Select this option to have all messages pertaining to this job saved to the Job Log.

Log Summary Only (the default selection)

Only general information about the job will be saved to the Job Log. Summary information includes the job source, destination, and session number, as well as any errors that may have occurred during job execution.

Log Disabled

Select this option when you do not want any messages about this job saved to the Job Log.

Alert options

You can use the Alert notification system to send messages about events in your restore operation. For more information on the Alert system, refer to Chapter 16, “Using the Alert Manager”.



To get to the Alert tab in the Options dialog box, click the Options button on the toolbar.

The screenshot shows the 'Alert' tab of a dialog box. It is divided into two main sections: 'Event' and 'Method'.
Under 'Event', there is a sub-section 'Send Alert when:' with five checkboxes:

- Job completed successfully
- Job incomplete
- Job canceled by user
- Job failed
- Virus detected

Under 'Method', there is a sub-section 'Send Alert via:' with five checkboxes:

- Broadcast
- Pager
- SNMP
- Trouble Ticket
- Email

At the bottom of the dialog, there is a checkbox labeled 'Attach job log' which is unchecked.

Event

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

Event	Description
Job completed successfully	All of the nodes and drives/shares were processed.
Job incomplete	Some nodes and/or drives/shares were missed.
Job cancelled by user	The user cancelled the job.
Job failed	The job was started but could not be completed.
Virus detected	Virus scanning does not apply for restore operations.


Method

Choose one or more of the following Alert notification methods. Before you choose a method, make sure that the Alert system is properly configured for this method. Refer to Chapter 16, “Using the Alert Manager” for more information.

Method	Description
Broadcast	Windows NT Broadcast's can be sent to specific users or groups.
Pager	Alphanumeric paging system.
SNMP	Simple Network Management Protocol (such as NetWare Management System (NMS) and HP Open View).
Trouble Ticket	Any print queue on your network. (Trouble Tickets can only be sent to non-Postscript printers.)
Email	Microsoft mail.

Attach job log

Select this option to include the job log information in the Alert message. (This option applies for Email only.)





COPYING FILES

The Copy function allows you to copy files and directories from one drive to another drive.

In this chapter, you will learn:

Page

- | | | |
|------|---|--|
| 6-4 | ➤ | How to use the Copy Manager |
| 6-5 | ➤ | How to perform a quick, basic copy job |
| 6-12 | ➤ | How to use advance copy options |
| 6-35 | ➤ | How to schedule a copy job |

Copy basics

The Copy feature allows you to copy (or move) files from one hard 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 tape. Or, you might want to copy files when your hard disk space is running low.

ARCserve's powerful copy capabilities are easy to use. With them, you can copy a drive, volume, directories, or files between:

- Your local Windows NT machine
- Remote Windows NT machines on your Microsoft network
- NetWare file servers

Copying files from remote Windows NT machines

ARCserve for Windows NT allows you to copy files from any Windows NT machine on your Microsoft network. Copying files from a Windows NT machine can be done in one of two ways:

- Copy an administrative shared drive
- Copy user shared files, directories, or a drive

Administrative shared drives

By default, Windows NT creates an administrative share for the root directory of each drive on the machine. Unless the administrative share is deleted, you can select it to be copied and ARCserve is able to copy all of the files and directories on the drive.

User shared drives and directories

In addition to creating default administrative shares, Windows NT allows you to designate specific files and directories as shareable. The owner of these files and directories has

specified that they are accessible to all. These files and directories are displayed in ARCserve's Browser, and with a click of the mouse button, you can select them to be copied.

Get familiar with the basics, then use advanced features and options

Once you get accustomed to the basics of setting up a copy job, explore ARCserve's copy options, filters, and scheduling techniques to really take advantage of ARCserve's power and flexibility.

Copy options are specifications that you can make to enhance and define each job. Jobs that you submit without setting copy options are jobs that accept the default copy option settings. To change the default copy options, refer to the 'Advanced copy features and options' section on page 6-9.

Filters allow you to be very explicit about the files and directories you want copied. Suppose you want to copy the majority of files and directories on a particular Windows NT drive. Instead of using the Browser to select each directory and file that you want, use the Browser to select the entire drive, then add filters to exclude the directories and files that you don't want included. For details about filters and how to use them, refer to Chapter 7, "Customizing Your Jobs".

Scheduling your copy job allows ARCserve to automatically re-submit jobs that you want to run on a regular basis. For details about scheduling your job, refer to the 'Scheduling your copy job' section on page 6-35.

Save important jobs to a script file

Saving your copy job settings to a script file makes it easier for you to resubmit jobs to the queue if and when you need to. For more information about scripts, refer to the 'Creating a script for your copy job' section on page 6-37.

Copy Manager The Copy Manager is shown below:

Run/schedule the copy job. Set options for the copy job. Add a filter to the copy job.

Select the Source tab to specify the source for the copy job and the Destination tab to specify the destination for the copy job.

The Browser area is where you specify the files and directories to copy and the disk to which they should be copied.

CPU Type	486
Number of CPU's	1
Total Memory	33,095,680
Free Memory	3,829,760

Host: TP_AS User: Administrator 12/20/94 2:49 PM

Copying files - the basic steps

Following are the basic steps for performing a copy job - that is, using the default options to copy files and directories from one hard disk to another. Advanced copy options are discussed later in this chapter.

Rebecca is working on a quarterly accounting report. Since her machine is not going to be backed up tonight, she wants to copy her critical files to another machine that is scheduled to be backed up.

Basic steps for copying files




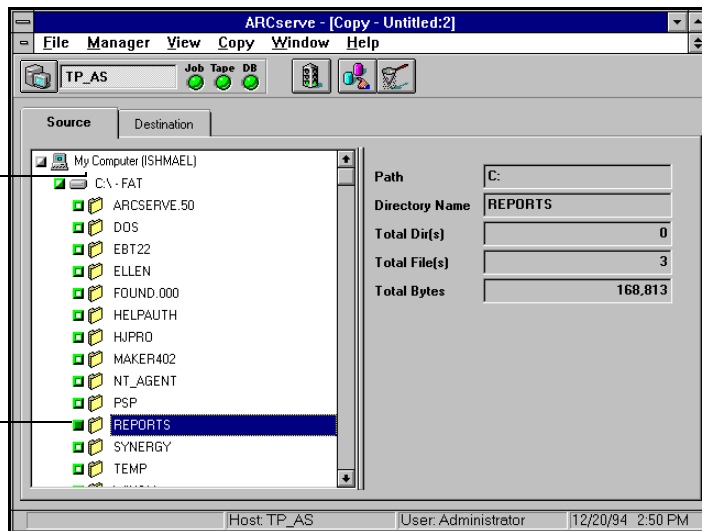
Using the above example, follow these directions to copy files:

1. Click the Copy button from the Quick Access window .
2. In the Source pane, select the files you want to copy .
You can only select one drive or volume, but you can select multiple directories or files on that drive or volume.

Rebecca wants to copy the entire "REPORTS" directory.

Rebecca double-clicks to expand the directory tree under "My Computer" until she reaches the "REPORTS" directory on her C drive...

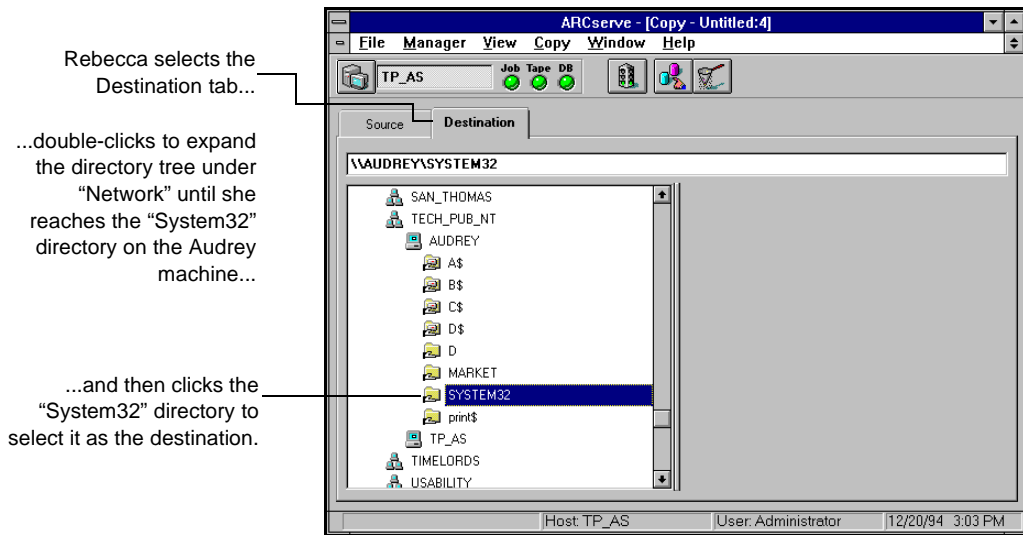
...and then clicks the box next to the "REPORTS" directory to select it as the source.




-
3. In the Destination pane, specify where you want to copy the files.

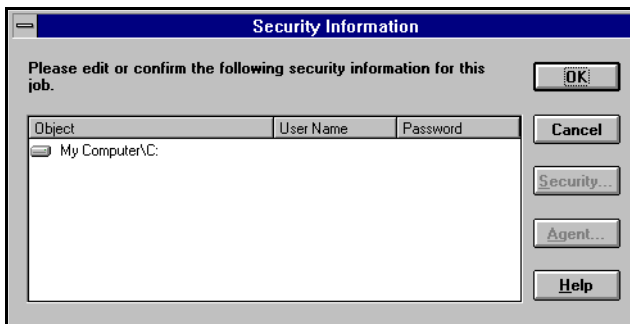
Again, refer to Chapter 3, “Basic Procedures”, for more information about using the Browser to select your destination.

Rebecca wants to copy the files to the root of the “System32” directory on the AUDREY machine.



4. Click the Run/Schedule button. 

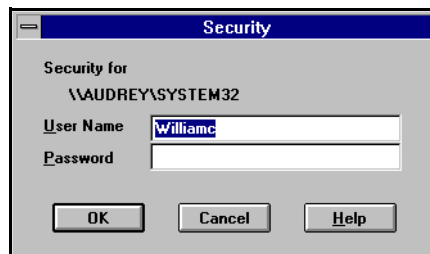
The Security Information dialog box opens:



This dialog box displays the security information for the source machine or share. Enter or verify the information that is shown. To enter or change the username or password, select the machine, then click the Security button.

5. Click OK.

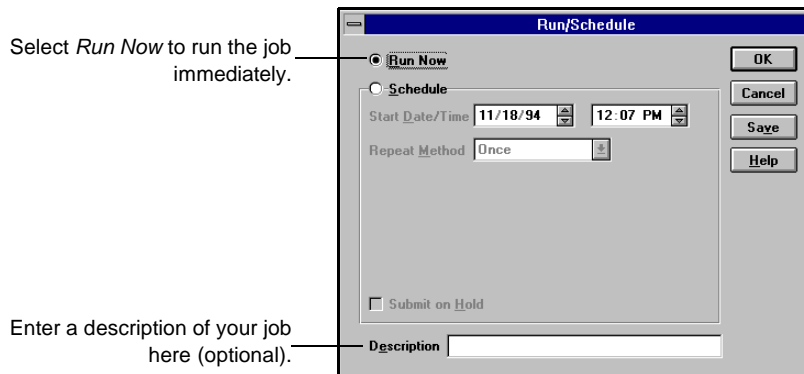
The Security dialog box opens:



Enter or verify the username and password for the destination machine.


6. Click OK to continue.

The Run/Schedule dialog box opens:



Select *Run Now* to run the job immediately.

Enter a description of your job here (optional).



The Run/Schedule dialog box enables you to run the job, or schedule the job for a later time.

Rebecca wants the job to run immediately. She confirms that Run Now is selected, then she clicks OK.

To save this job to a script file, click the Save button in this dialog box.

7. Click OK to submit this job.

Advanced copy features and options

Now that you know how to perform a basic copy job, you are ready for ARCserve's advanced copy features. These include using the hidden, right-click menu to provide security information, applying options to your job, scheduling your job, and creating job scripts.

Using the hidden menu to provide security information

ARCserve's Copy Manager has a hidden menu in the Browser that allows you (if necessary) to provide security information so that you can access remote machines. Refer to page 6-10 for more information.

Applying options to your job

There are a number of advanced copy options you can add to your copy job. These options ensure that you get exactly the copy job that you want. Refer to page 6-12 for more information about the options that are available.

Scheduling your job

ARCserve allows you to schedule your copy job so that it runs at regular intervals. Refer to page 6-35 for more information.

Creating job script

ARCserve allows you to save all of your copy job settings into a script that can be used again at a later time. Refer to page 6-37 for more information.



Filtering your copy job is discussed in Chapter 7, "Customizing Your Jobs".

Providing security information

You must provide a valid username and password for the machines or shares that you wish to back up. You can enter the username and password:


- By clicking on the machine or share with your *right* mouse button.
- When attempting to expand the machine/share to view directories. (You will be prompted for the information.)
- When running a backup job. (You will be prompted for the information.)

Browsing

The Browser uses your local username and password by default in order to access machines or shares for browsing. If your local username and password are not valid on the remote machine or share, then you will be prompted to enter the information when you attempt to expand that node.

You can also enter a username and password from the Browser as follows:

1. Select the machine with your right mouse button.
A popup menu is displayed.
2. Select Security from the pop-up menu.

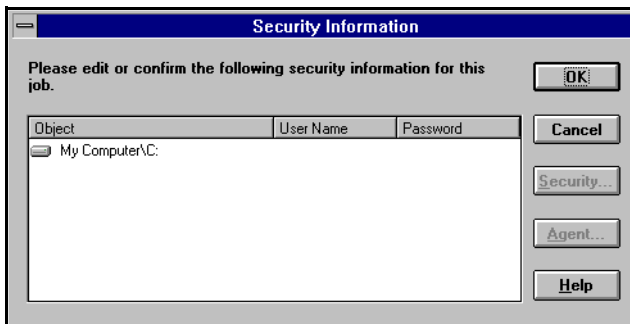


3. Enter a valid username and password.
4. Click OK.

Copying

You must also enter or verify the username and password for a machine each time you submit a job. If you have already provided account information for a machine, either in the Browse mode or in a previous copy or backup job, ARCserve will retain this information and simply ask you to verify it. (This information is stored for future sessions.)

When you submit a copy job, the Security Information dialog box opens:



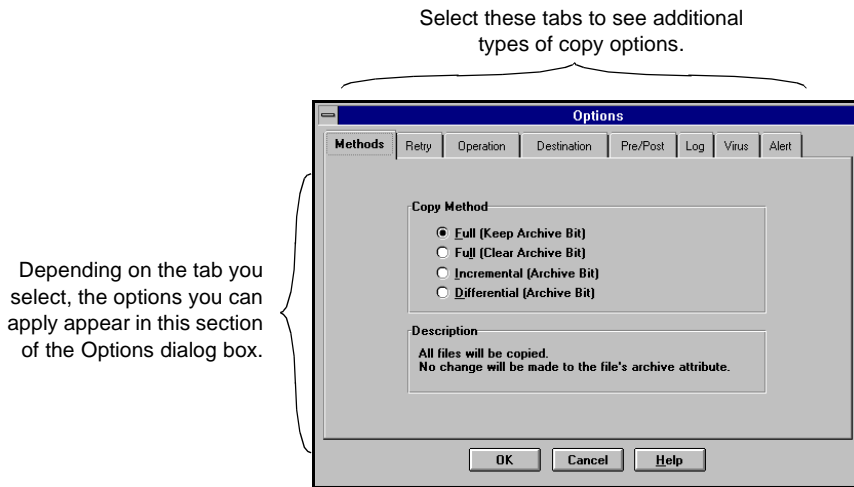
This dialog box displays the username and password for the source machine. To change this information, select the machine, then select the Security button.

Applying advanced copy options

Once you know how to perform a basic copy job, you can customize your jobs by applying advanced options. To apply advanced options to your copy jobs, follow the steps below:

1. From the Copy Manager, click the Options button. 

The Copy Options dialog box appears:



2. Select the tab that contains the options you want.

For more information about each advanced option:

Option	Refer to:
Methods	page 6-14
Retry	page 6-19
Operation	page 6-23
Destination	page 6-25
Pre/Post	page 6-30
Log	page 6-31

Option	Refer to:
Virus	page 6-32
Alert	page 6-33

Methods

Select the Methods tab in the Options dialog box to select a Copy Method. Copy methods determine what files will be copied and, after each file is copied, what ARCserve will do with the file's archive bit.



The following table briefly outlines each option:

Method	What is copied?	Archive bit reset after copy job?
Full (Keep Archive Bit)	All source files.	No.
Full (Clear Archive Bit)	All source files.	Yes.
Incremental (Archive Bit)	Source files whose archive bit is turned on	Yes.
Differential (Archive Bit)	Source files whose archive bit is turned on	No.

What is an archive bit?

An archive bit is a marker that indicates whether or not a file has been modified. Every file has an archive bit. This bit is set (turned on) when the file is created or modified. Therefore, if a file's archive bit is set, that file has changed since it was last copied (or backed up).

Why turn off a file's archive bit?

Turning off the file's archive bit is a procedure that affects what ARCserve will copy the next time. An archive bit that is off

indicates to ARCserve that this file has not changed since it was last copied.



There are other ways a file's archive bit can be turned on and off:

- Setting/Resetting the file's Archive attribute through another application.
- Running ARCserve backup jobs with a backup method that resets the archive bit.

Be careful when selecting a copy method that relies on the archive bit. There may be times when a file is not copied because another application has changed its archive bit setting.

The copy method you select determines whether ARCserve will copy all files, or only files whose archive bit is on. It also determines whether ARCserve will turn off the archive bit after the file has been copied.

Full (Keep
Archive Bit)

This is the default copy method. This method copies all of the files that you selected in the Source Browser. After a file is copied, its archive bit is not changed, meaning, the archive bit remains at whatever setting it was at before the copy job.

Full (Clear
Archive Bit)

This method copies all of the files that you selected in the Source Browser. After a file is copied, ARCserve resets (turns off) its archive bit.

Incremental
(Archive Bit)

This method copies only those files that you selected in the Source Browser **THAT HAVE THEIR ARCHIVE BIT SET**. After a file is copied, ARCserve resets (turns off) its archive bit.

Differential
(Archive Bit)

This method copies only those files that you selected in the Source Browser **THAT HAVE THEIR ARCHIVE BIT SET**.

After a file is copied, ARCserve DOES NOT reset its archive bit. Therefore, the archive bit remains set.

Advantages and disadvantages of each method

Depending on your needs, each method has advantages and disadvantages. For example, a full copy guarantees that all of your data is current and is in one place. But, full copy jobs take longer to perform than incremental or differential jobs. If you have a large amount of data to be copied, or limited disk space on the machine where you are copying your data to, it may be unreasonable to perform full copy jobs every day.

Deciding which method is right for you

Some of the questions you should be asking yourself when deciding upon a copy strategy are:

- How much data gets copied?
- What percentage of my data changes each day?
- How long do I want the copy job to take?
- What type of drive am I copying files to?

In general, if you have large amounts of data to be copied every day, differential or incremental copy jobs may be your best solution.

Differential copy jobs will only include files that have changed since the last full copy job.

Incremental copy jobs will only include files that have changed since the last full or incremental copy job, and in general are considered the quickest method for copying files.

Therefore, if you are copying files to a hard drive on another machine, you would probably want to schedule one full copy job and then use incremental jobs for subsequent days. This ensures that all of your data is current and reduces the amount of time it takes to perform a copy job.

Considerations
for WORM drives

A WORM (Write Once Read Many) drive is a storage device that uses an optical medium that can be recorded only once. You can map a WORM device like you would map any network drive. But, because it is a “write once” device, if you schedule a full copy of your C drive to your WORM device each day, files will not overwrite older copies of themselves on the drive (the way they would if you copied a file from a diskette to your hard drive). Instead, you will have separate full copies of your C drive (similar to backing up files and appending the data to a tape).

Therefore, your copy strategy when using WORM devices might be to perform a full copy job once and then use either incremental or differential jobs on subsequent days.

If you are having trouble deciding whether to use incremental or differential methods, refer to the following copy examples for WORM drives:

Scenario #1: Using
incremental copy jobs

You perform a full copy on Monday and then incremental copy jobs on Tuesday through Friday. Saturday, you need to recover your entire C drive. You have to copy Monday’s full copy job and Tuesday through Thursday’s incremental copy jobs to recover all the files for your C drive.

The advantage in this example is the speed of the incremental copy jobs performed Tuesday through Friday. The disadvantage is the increased amount of time it takes to recover files in the event of an emergency.


Scenario #2: Using differential copy jobs

You perform a full copy on Monday and then differential copy jobs on Tuesday through Friday. Saturday, you need to recover your entire C drive. You have to copy Monday's full copy job and Friday's differential copy job to recover all the files for your C drive.


The advantage in this example is the short amount of time it takes to recover files in the event of an emergency. The disadvantage can be the increased amount of time it takes to perform differential copy jobs (vs. performing an incremental copy job).

Retry options

Open File Retry



Select the Retry tab in the Options dialog box to select the action ARCserve should take if it comes across an open file during the copy job.



Retry Immediately

If you select *Retry Immediately*, ARCserve will try to copy the file again immediately (based on the numbers in the *Number of Retries* and *Retry Interval* fields). If the file is still unavailable, this information is written to ARCserve’s Log, and the job is labeled “Incomplete”.

Retry After Job

If you select *Retry After Job*, ARCserve will try to copy the file again after the rest of the job has been completed (based on the numbers in the *Number of Retries* and *Retry Interval* fields). If the file is still unavailable, this information is written to ARCserve’s Log, and the job is labeled “Incomplete”.



You don’t have to limit yourself to one selection here. You can select neither, one, or both of these options (*Retry Immediately* and *Retry After Job*).

Number of Retries

Indicate how many times ARCserve should try to copy a file, if the attempt is unsuccessful because the file is inaccessible.


Retry Interval

Indicate how long (in seconds) ARCserve should wait between retries. (ARCserve will not copy other files during this time.)


File Sharing

Select the Retry tab in the Options dialog box to select a File Sharing option.

When ARCserve attempts to open a file to copy it, it needs to know whether you want subsequent users and processes to be able to access this file and, if they can, whether they can save data to it. If the file is already open by another process and ARCserve cannot open the file with the selection you make here, ARCserve will then use the *Open File Retry* option that you specified to attempt to copy the file again. The default is *Use Deny None if Deny Write fails*.




Use Deny None if Deny Write Fails



ARCserve will attempt to open the file in *Deny Write* mode first (which allows users to read from the file but not write to it). If the file is already in use, ARCserve will place the file in *Deny None* mode (which lets users read from and write to the file) and then copy it.

Use Lock Mode if Deny Write Fails



ARCserve will attempt to open the file in *Deny Write* mode first (which allows users to read from the file but not write to it). If the file is already in use, ARCserve will lock the entire file, and then copy it.

Lock Mode will ensure that the most current version of the file is copied.

Deny Write



Deny Write means that while ARCserve has the file open and is copying the file to the destination, another process can only open the file for reading. A file can only be opened in *Deny Write* mode if it is not already open by another process.

This mode is the safest for copying, as no other process can write to the file during the copy job. If, however, another process has the file open before ARCserve gets a chance to open it, ARCserve will not be able to open the file, and the file will not be copied to the destination (unless you specified an *Open File Retry* option which was discussed in the previous section).

Deny None



Deny None mode will allow another process to read or write to the file, regardless of whether ARCserve opens it first or opens it after another process already has it open. This method ensures that your files are up-to-date, although your copy job may not be.

Deciding which option to select

If you use applications (such as 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 copy will be consistent.

If you want to ensure that only the most current version of every file is copied, you should select a *Deny Write* or *Lock Mode* option.

If you are still having trouble deciding which option to select, refer to the following examples:

Scenario #1: An online database being used 24 hours a day

Your customer service department is open 24 hours a day and the staff uses a database application to record customers' orders. The database is copied every night. Since you do not want customers to wait while the database is copied, you select the Deny None method so that the copy job will not interfere with the order taking.

Scenario #2: Hardware problems and a tight deadline

Your sales department is working long hours to finish a presentation for an important potential client. Since the machine that contains the files has been having hardware problems lately, you want to make sure the files are copied to another machine regularly. You select the Use Lock Mode if Deny Write Fails method so that you can have the most current copies of the presentation files available.

Operation options

Operation

Select the Operation tab in the Options dialog box to select any options you want ARCserve to use. By default, none of these options are initially selected.



Mirror

This option creates a mirrored version (an exact duplicate) of the source.



Be careful when selecting this option. In order to create an exact duplicate, ARCserve DELETES any files on the destination that are not on the source and then ARCserve copies all of the files from the source to the destination. As a result, your source and destination are always the same.

If you want to mirror a particular directory (for example, one you use constantly and must always have access to), you should use the *Mirror* option in conjunction with the *Automatic Repeat Interval* option (set when scheduling the job). This way you will always have a current mirrored directory.

For example, if you specify an *Automatic Repeat Interval* of 15 minutes, and specify in the Destination Options dialog box to *Overwrite with Newer Files Only*, all files that have changed during the previous 15 minutes will be copied to the destination directory. In this way, you will have a truly accurate mirrored directory every 15 minutes.

Delete Files


Select this option if you want to delete the source files after they are successfully copied to the destination. Using this option allows you to move files from one location to another.

Create Empty Directories

Select this option if you want ARCserve to create directories on the destination even if the source files or directories are empty, or, if after filtering, the directories will be empty.

Database

Select the Operation tab in the Options dialog box to indicate whether or not ARCserve should keep copy information in its database. The default is *Record Job Information Only*.



Record Job Information Only

Select this option if you want a record of this copy job logged in the ARCserve database. Session and file information are not kept in the database for copy jobs, only general job information is maintained.


Disable Database Recording

Select this option if you do not want this job recorded in the ARCserve database. You may want to do this when you have limited disk space.

Destination options

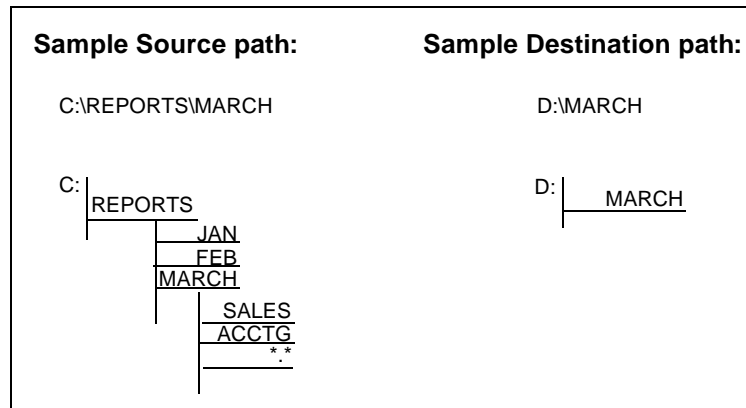
Directory Structure

Select the Destination tab in the Options dialog box to select the method ARCserve should use to create directories on the destination that you select. The default is *Do Not Create the Base Directories*.



Suppose you want to copy the path C:\REPORTS\MARCH. There are several files directly under the MARCH directory, and there are two subdirectories, \SALES and \ACCTG. You want to copy to the MARCH directory on your D drive.

Source path and destination path for Directory Structure options examples




Do Not Create the Base Directories

A “base directory” is considered the last directory selected in the source path. This option will not create the base directory on the destination path, but will create all subdirectories that may reside below the source base directory.

Using this option to copy C:\REPORTS\MARCH to D:\MARCH would produce the following:

Since MARCH is the base directory, it is not created on the destination. However, all files and subdirectories below it are created.




Create Directories from the Base

This option will create the destination path beginning from the base directory.

Using this option to copy C:\REPORTS\MARCH to D:\MARCH would produce the following:

Using the above example, MARCH is the base directory. As a result, everything on the source beginning with the MARCH directory is created on the destination.




Create Entire Path from the Root

This option will create the entire source path (except the root drive or volume name) on the destination. No files from any parent directories are copied. Only the directory path to the base directory is created on the destination.


Using this option to copy C:\REPORTS\MARCH to D:\MARCH would produce the following:

Using the above example, C: is the root. As a result, the path to the MARCH directory is built beginning from the REPORTS directory. Note that no files from parent directories are copied - only the path to the base directory is created.



File Conflict Resolution

Select the Destination tab in the Options dialog box to 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. The default is *Overwrite All Files*.



Overwrite All Files

This option will copy all source files to the destination regardless of conflicting file names. Existing files on the destination will be overwritten by the files from the source.

Rename Files

This option copies 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. Filenames without extensions will be renamed with extensions “._0”, “._1”, “._2”, etc.

For NTFS and HPFS long names, filenames with or without an extension will be renamed with extensions “._0”, “._1”, “._2”, etc.

Skip Existing Files

If you select this option ARCserve will not copy a source file if a file with the same name already exists on the destination.

Overwrite with Newer Files Only

If you select this option ARCserve will only copy source files whose modification date is later than the modification date of the file with the same name on the destination. Source files whose modification date is earlier will not be copied to the destination.

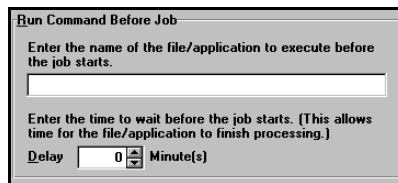
*Confirm
Overwrites*

This option is only applicable when running an immediate (Run Now) job. Before ARCserve tries to copy the source file, you will be prompted to confirm that you want the file on the destination to be overwritten.

Pre/Post options

Select the Pre/Post tab in the Options dialog box to run a command on your ARCserve machine before and/or after the copy job is executed.

The types of commands you can run include batch files and executable files that can be run from a command line. You should not try to run a batch file or application that requires a user to react to a screen prompt.



Run Command Before Job

Enter the name of the file/application to execute before the job starts.

Enter the time to wait before the job starts. (This allows time for the file/application to finish processing.)

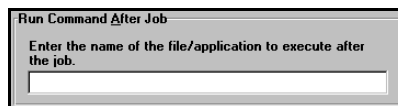
Delay 0 Minute(s)

Run Command Before Job

Enter the name of the file or application to be executed on your ARCserve machine before the job takes off.

Delay

Enter the time to wait (in minutes), before the copy job will start. This period of time gives the *Run Command Before Job* application or file time to process.



Run Command After Job


Enter the name of the file/application to execute after the job.

Run Command After Job

Enter the name of the file to be executed on your ARCserve machine after the job has completed.

Log options

Select the Log tab in the Options dialog box to determine the level of detail you will see in the Job Queue Log or Job View Database Log.



Log All Activity

If you select this option, all of the activity that occurred while the job is running will appear in the log.

Log Summary Only

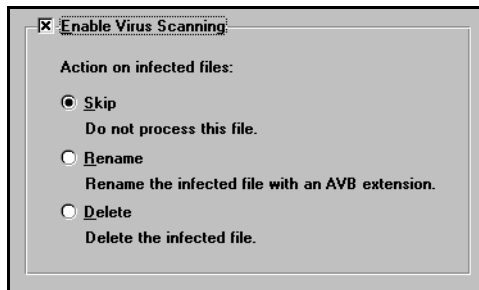
If you select this option, you will get a summary of the job (which includes source, destination, session number, and totals) along with any errors that may have occurred.

Log Disabled

If you select this option you will not get any information in the log.

Virus scanning options

Select the Virus tab in the Options dialog box to enable virus scanning.



Enable Virus Scanning

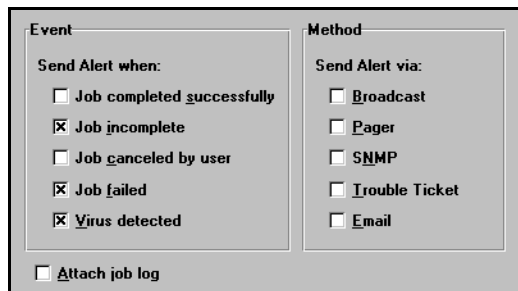
Select this option to scan files for viruses before they are copied.

The following table describes the virus scanning options:

Option	Description
Skip	Do not copy the infected file.
Rename	Rename the infected files with the extension AVB. If a file with the same name and an AVB extension already exists, then the infected file is renamed with the extension AV0. The number is incremented with each subsequent occurrence (AV0, AV1, etc.)
Delete	Delete the infected file.

Alert options

Select the Alert tab in the Options dialog box to configure the Alert notification system to send messages about events in your copy operation. For more information on the Alert system, refer to Chapter 16, “Using the Alert Manager”.



Event

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

Event	Description
Job completed successfully	All of the nodes and drives/shares were processed.
Job incomplete	Some nodes and/or drives/shares were missed.
Job cancelled by user	The user cancelled the job.
Job failed	The job was started but could not be completed.
Virus detected	A virus was detected. Refer to the 'Virus scanning options' section on page 6-32 for more information.

Method

Choose one or more of the following Alert notification methods. Before you choose a method, make sure that the

Alert system is properly configured for this method. Refer to Chapter 16, "Using the Alert Manager" for more information.

Method	Description
Broadcast	Windows NT Broadcast's can be sent to specific users or groups.
Pager	Alphanumeric paging system.
SNMP	Simple Network Management Protocol (such as NetWare Management System (NMS) and HP Open View).
Trouble Ticket	Any print queue on your network. (Trouble Tickets can only be sent to non-Postscript printers.)
Email	Microsoft mail.


Attach job log

Select this option to include the job log information in the Alert message. (This option applies for Email only.)

Scheduling your copy job

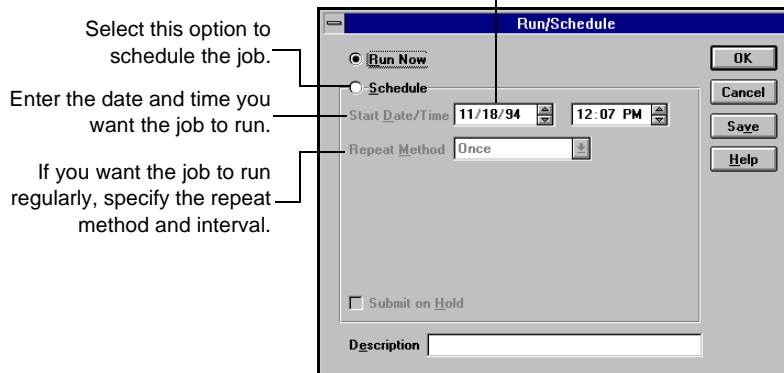
With the source and destination specified, you can schedule the job to run.

The basic procedure for scheduling your copy job is discussed below. Refer to Chapter 7, “Customizing Your Jobs” for detailed instructions about job scheduling.

1. Click the Run/Schedule button. 


After entering the security information for your source and destination machines, the Run/Schedule dialog box opens:

Double-click to display a calendar to help you select the date for the job to run.



2. Select Schedule.
3. Specify the date and time you want the job to run.

The current date and time appear initially. Specify the time and date using one of the following methods:

- Type the specific date and/or time manually.
- Use the spin buttons  to select the date and time.

-
- Double-click in the date field to display a calendar to select a specific date.



4. (Optional) Specify a repeat method.

If you want this job to run regularly, you need to provide the information in the *Repeat Method* fields.


Creating a script for your copy job

You can save your copy job as a script. The script will contain the source, destination, and options that you selected, as well as the schedule information. It will also contain any filters you created to include and exclude files and directories.

The basic procedure for creating a script is discussed below. Refer to Chapter 7, “Customizing Your Jobs” for detailed instructions about creating and using scripts.

1. Set up the copy job.

Select the source, destination, and any options or filters.

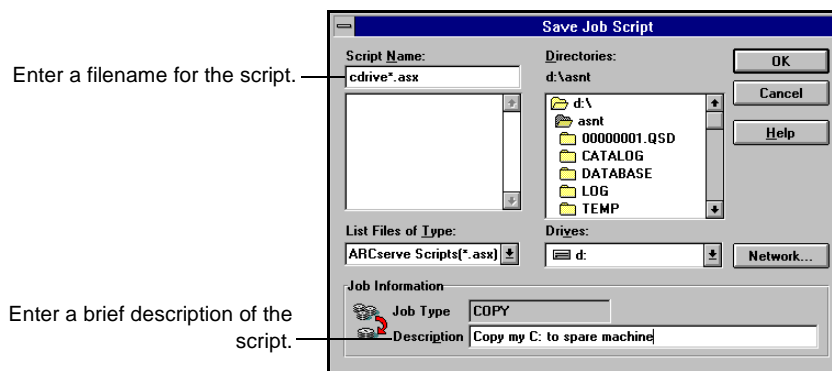
2. Click the Run/Schedule button. 

After entering the security information for your source and destination machines, the Run/Schedule dialog box opens.


3. From the Run/Schedule dialog box, click the Save button.

You can also save a script by selecting *Save As* or *Save* from the File menu.

The Save Job Script dialog box appears.



4. Give the script a name.



By default, ARCserve scripts can be stored in the ARCserve home directory with a default extension “.asx”. You can, however, specify a different directory in which to save your script.

5. (Optional) Enter a description of the job .
This can help you identify the correct script at a later date.
6. Click OK.



CUSTOMIZING YOUR JOBS

This chapter discusses filters that you can set for your jobs. It also discusses how to schedule jobs and create scripts.

In this chapter, you will learn:

<i>Page</i>	
7-2 >	What types of filters are available
7-5 >	How to apply a filter
7-15 >	How to modify or delete a filter
7-17 >	How to schedule a job
7-28 >	How to create scripts for your job

Filter basics

About filtering

Filters allow you to include or exclude files and directories from your backup, copy, or restore jobs, as well as from the utilities, such as count and purge. For example, you could create a filter to back up only those files that were modified on or after a certain date. You could also create a filter to exclude a particular directory from being restored. Use the filters to focus on the files you want.

For backup and count jobs, filtering can be performed on a per machine basis. This means you can include a directory from one machine and exclude the same directory from another machine. These filters make it easy for you to create a job that best suits your needs, one that contains only the data that you want.

Exclude over include

One thing you should be aware of is that exclusions will always take precedence over inclusions. For example, if you add a File Pattern filter to include only files that have an *.exe* extension, and you add a Directory filter to exclude your *\system32* directory, all *.exe* files in the *\system32* directory will be excluded.

Include only

Specifying to “Include” files will result in a backup or restore containing only those files which satisfy the filter specifications. For example, suppose in the source area you selected to back up your entire local hard drive, and you then set up a Directory filter to include files in the *\system32* directory. The result would be that ARCserve would only back up files from your *\system32* directory. No other files would be backed up.

Default filter settings - everything is included in the job

If you don't specify any filters for a job, ARCserve will include all files that are on the selected source.

Available 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 created, modified, or last accessed before, on, or after a specific date or within a date range

Below is a brief description of the filters that you can use.

File Pattern filter

The File Pattern filter allows you to select specific file names to include or exclude from your job. If you don't know the specific file name, you can provide as much of the file name as you know, then use wildcards to fill in the blanks.

Directory filter

The Directory filter allows you to exclude or include files in entire directories. As with the File Pattern filter, you can enter an entire directory name, or provide a pattern that the directory name follows. For example, setting a filter to include the *\user* directory will include all of the files in the *\user* directory with your job. If you enter *use**, all directories that begin with the letters *use* will be included with your job (*user*, *users*, *userdoc*, etc.).

File Attributes filter

The File Attributes filter allows you to filter files based on whether they are hidden, system, archived, read only, or any combination of the four.



File Modified/
Created/
Accessed filters

These filters allow you to filter files based on a threshold date or date range. Specify a date or date range and whether to include or exclude files that were created, modified, or last accessed before, after, or between the date range.

Applying filters - the basic steps

Machine-level filters

Machine-level filters can be created for backup and count jobs. These filters apply to one specific machine, not the entire job. (A backup or count job can have machine-level and job-level filters for the same job.) If you want to add a filter that applies to the whole job, refer to the 'Job-level filters' section below.

To apply a machine-level filter:

1. Highlight the machine you want to filter .
2. Click your right mouse button .



3. Select Filter.

The Filters window appears. In it, you will create your machine-level filters. Refer to the 'Job-level filters' section below for information about the Filters window.

Job-level filters

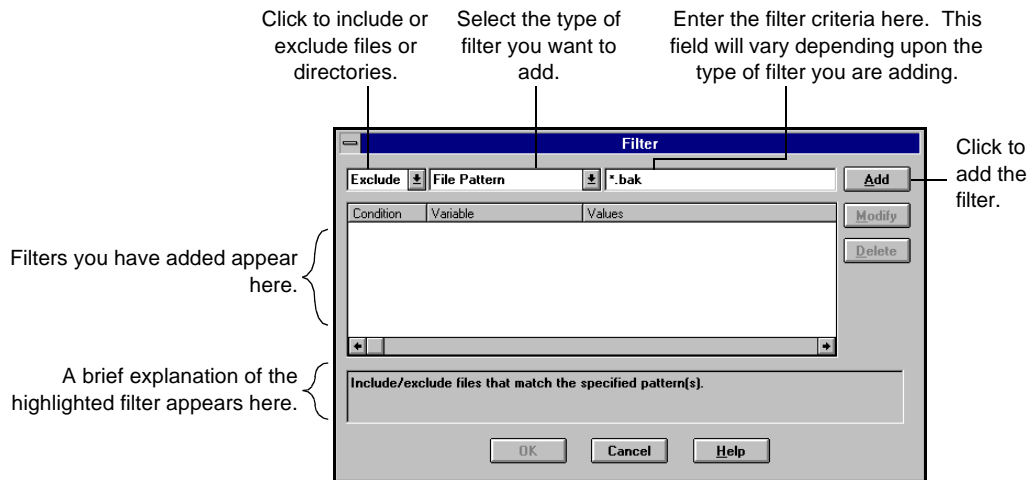
Job-level filters can be created for backup, copy, and restore jobs, as well as for the utilities, such as count. Job-level filters are applied to the whole job, not to a specific machine.

To apply a job-level filter:

1. Click the Filter button. 



The Filters window appears. In it, you will create all your job-level filters.



2. Decide whether to include or exclude the files or directories.

If you select *Include*, only files or directories that match the criteria you specify will be backed up, copied, restored, counted, or purged.

If you select *Exclude*, all files will be backed up, copied, restored, counted, or purged EXCEPT for those that match the criteria you specify.


3. Specify the filtering information.

Select the type of filter you want to add and enter the criteria (file pattern, directory, date(s), or file attributes) for that filter.

Refer to the sections that follow in this chapter for information about the various filters.

4. Click the Add button.

File Pattern filter

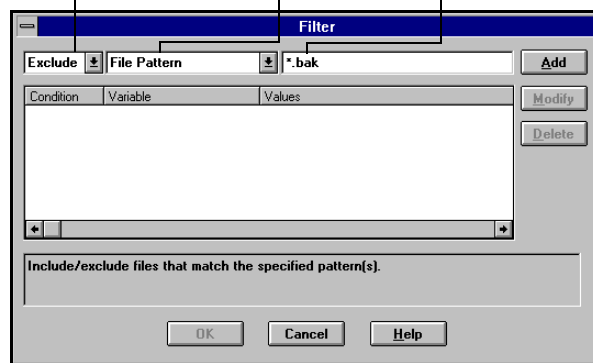


Use a File Pattern filter to include or exclude certain files from a job. You can specify a particular file name or you can use wildcards to specify a file pattern.

You are copying files from your reports directory to another machine. In order to save space on that machine, you want to exclude all files with a .bak extension (these are the backup copies of the reports).

To accomplish this, you would do the following:

Select *Exclude* to exclude all files that have a *.bak extension.




Entering file names and patterns

There are several methods for entering the file name or pattern to filter. The table on the following page shows you how to specify path and file names for your computer and remote machines.



You cannot enter *.* as your file pattern to include files.



Windows NT
machine/
NetWare server
file filtering

While reading the following table, keep in mind that:

- DRIVE: represents a Windows NT machine drive.
- PATH represents a series of directories separated by backslashes (\).
- FILE is the name of the file to filter. For machines using the File Allocation Table (FAT) file system, the file name must be in the 8.3 format. For machines using the Windows NT File System (NTFS), NT 3.5 FAT, or the High Performance File System (HPFS), the file name can contain up to 256 characters.

Pattern	What it means	Example
DRIVE:\PATH\FILE	This is an absolute path on your local Windows NT machine drive.	C:\SYSTEM*.FOT C:\BATAUTOEXEC.BAT
\\MACHINE\PATH\FILE	This is an absolute path on a remote machine.	\\ENG\DOS\SYSTEM.BAT
\PATH\FILE	This is an absolute path on any Windows NT machine drive.	WINNT\SYSTEM*.FOT \BATAUTOEXEC.BAT
DRIVE:FILE	The file may be in any directory on the specified Windows NT machine drive.	C:*.FOT C:AUTOEXEC.BAT
FILE	The file may be in any directory on any drive.	*.FOT AUTOEXEC.BAT

Directory filter

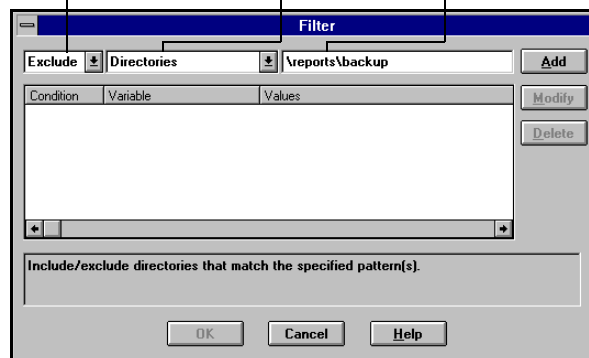


Use a Directory filter to include or exclude specific directories from a job. All files under the directory that you specify will be included or excluded, depending on your choice. You can enter an entire directory name or provide a pattern that the directory name follows.

You are copying files from your *reports* directory to another machine. In order to save space on that machine, you want to exclude all files from the *\reports\backup* directory (this is the directory that contains the backup copies of the reports).


To accomplish this, you would do the following:

Select *Exclude* to exclude the directory *\reports\backup*.



Entering directory names and patterns

There are several methods for entering the directory name or pattern to filter. The table on the following page shows you how to specify paths for Windows NT machines and NetWare servers.




Windows NT
machine/
NetWare server
directory filtering

While reading the following table, keep in mind that:

- DRIVE: represents a Windows NT machine drive.
- PATH represents a series of directories separated by backslashes (\).
- DIR_NAME is the name of the directory to filter.

Pattern	What it means	Example
DRIVE:\PATH	This is an absolute path on your local Windows NT machine hard drive.	C:\WINNT\MSAPPS C:\WIN*
\\MACHINE\PATH	This is an absolute path on a remote machine.	\\ENG\DOS
\PATH	This is an absolute path on any Windows NT machine.	\WINNT\MSAPPS \WIN*
DRIVE:DIR_NAME	The directory may be in any other directory on the specified Windows NT machine.	C:\MSAPPS C:\MS*
DIR_NAME	The directory may be in any other directory on any Windows NT machine.	MSAPPS *APPS

File Attributes filter

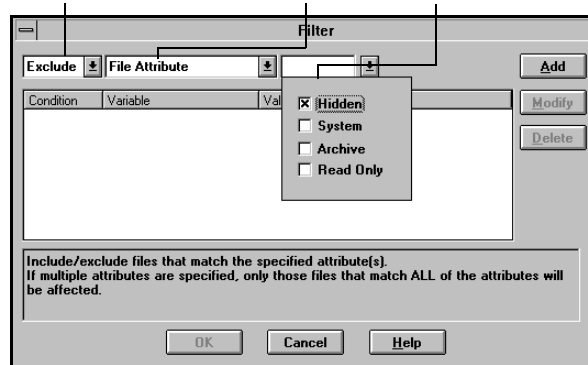


Use the File Attributes filter to include or exclude specific types of files from a job.

You want to back up your entire machine, but you want to exclude hidden files.

To accomplish this, you would do the following:

Select *Exclude* to exclude files with the attribute *Hidden*.



Types of file attributes

There are four types of file attributes you can choose from. Select as many file types as you want.

- **Hidden** - These are files that you cannot normally see when you get a directory listing. Examples of hidden files are BOOTSECT.DOS and IO.SYS.
- **System** - Files of this type are unique to the machine you are using. Examples of system files are BOOT.INI and MSDOS.SYS.
- **Read Only** - Files marked read-only are files that cannot be modified.
- **Archive** - Files whose archive bit is set.



ARCserve will only back up or restore the files that **EXACTLY** match the requirements you specify. For example, if you want to run a backup job that includes hidden files and system files, you would have to add two File Attributes filters to the job: One with Include and System Files selected and another with Include and Hidden Files selected. If you placed only one File Attributes filter on the job with Include, Hidden Files, and System Files all selected, ARCserve would only back up files that have 'hidden' **and** 'system' attributes, not hidden files and system files.

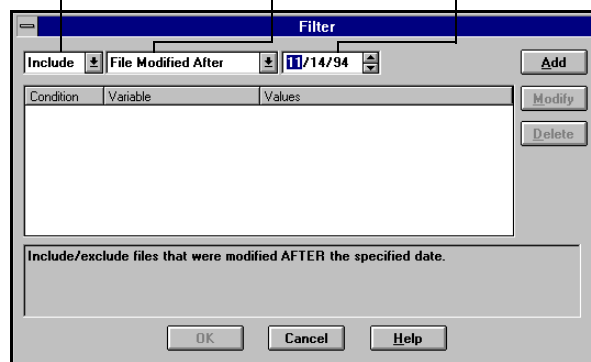
File Modified/Created/Accessed filter s

Use the File Modified/Created/Accessed filters to include or exclude files based on a specific date and time or a date range.

You backed up your entire machine last night. You want to back up your machine before you leave today. However, in order to save time, you only want to include files that have been modified today.

To accomplish this, you would do the following:

Select *Include* to include files modified on or after today.



Filter types

There are three types of filters from which to choose:

- *Modified* - will include or exclude files last modified (changed) on this date.
- *Created* - will include or exclude files created on this date.
- *Accessed* - will include or exclude files last accessed (used but not changed) on this date.



The *Created* and *Accessed* filter types are only available for NTFS (NT File System) and HPFS (High Performance File System).

Filter criteria

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* - You must specify two dates for this selection. Files whose date falls between the two dates will be included or excluded from the job.
- *Within* - Selecting this option requires you to enter a number of days, months, or years to be used for the comparison (for example, *within* the past 12 months). Enter a number between 1 and 250, inclusive. ARCserve will use the current date as the basis from which to count back.

Specifying dates

You can specify a date in two ways:

- Enter a specific date.
- Double-click in the date field to display a calendar to select a specific date.

Click to see the prior month.

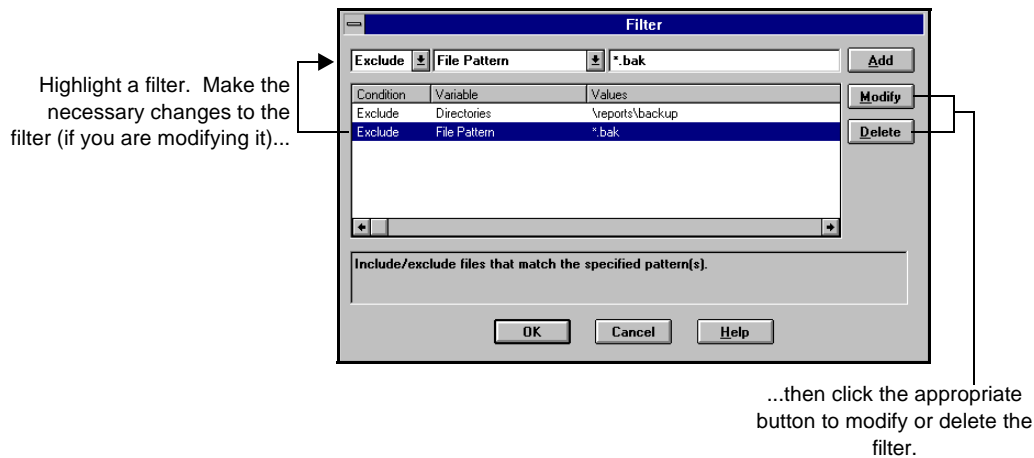


Click to see the next month.

Highlight the date you want and click to select it.

Viewing, modifying, or deleting existing filters

You can view the filters you've created by displaying the Filters window. You can also change or delete a filter from within this window.



Modifying existing filters

To modify a filter that has already been added to a job:

1. Highlight the filter to modify.
The filter will move up into the entry section of the window, complete with the information that you originally entered.
2. Make your changes to the filter.
You can change the values in any field.
3. Click the Modify button.
4. Click OK when done.

Deleting existing filters

To delete a filter:

1. Highlight the filter to delete.
2. Click the Delete button.

The filter will be deleted from the list of applied filters for the job.

Scheduling your job - the basic steps

When you have finished specifying your job requirements, you are ready to submit the job to the queue. This section describes how to schedule your job to run immediately or at a later date and time.



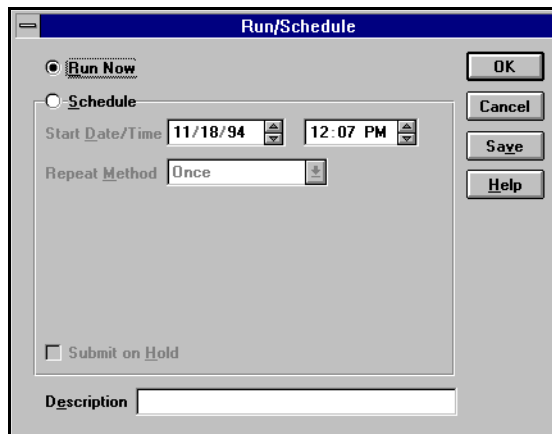
Before you can specify scheduling information, you must have specified the source and, if required, the destination for your job.

Follow the steps below to schedule your jobs:

1. Click the Run/Schedule button.



This button appears on the toolbar in most ARCserve Manager windows. When you click it, the Run/Schedule dialog box appears, as shown below.



2. Select when you want the job to run .

If you want the job to run immediately, select the *Run Now* option.

If you want the job to run later and/or you want to specify a repeat interval for the job, select the *Schedule* option.

Running a job immediately

This section describes how you can send your job directly to the queue to run it immediately.

When **should** you select *Run Now*?

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.

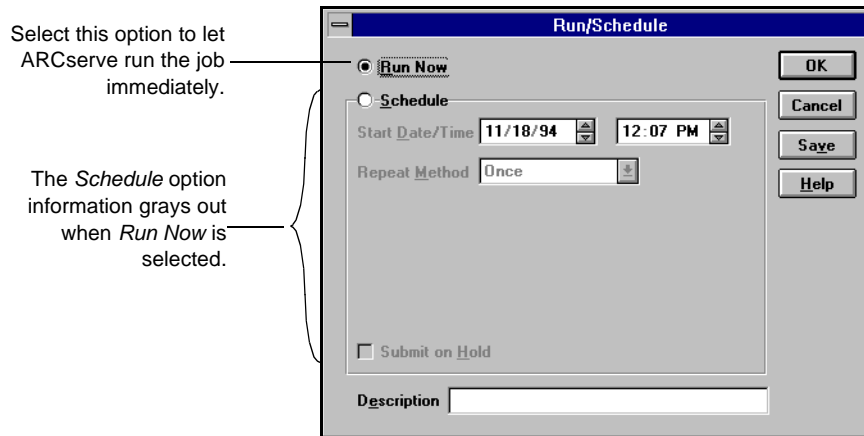
When **shouldn't** you select *Run Now*?

If your tape drive is currently busy and you want to run a backup job as soon as the drive is free, you should not select *Run Now*. Instead, you should schedule your backup job, keeping the current date and time. This way, when ARCserve discovers that the tape drive is busy, it will automatically retry the backup job until the drive becomes free.

If you select *Run Now*, ARCserve will report that the tape drive is busy and the backup job will not be submitted to the job queue. Instead, you will be asked to wait until the drive is available before you can resubmit the job.

To run your job immediately:

1. In the Run/Schedule dialog box, select Run Now and click OK.



You will be taken to the Job Queue Manager so that you can watch the progress of your job.

Scheduling your job

If you want your job to run at a specified date and/or time, and/or you want the job to repeat regularly (automatically resubmitted to the job queue), you need to specify this information with the Schedule option.

There are two aspects to specifying scheduling information:

- Specify a date and time - this is the least you have to do.
- Specify a repeat method - this is optional. Do this only if you want ARCserve to resubmit the job to the queue after it has finished. You may want to set up backup jobs with repeat intervals to implement a tape rotation scheme for your network.

When should you schedule your job?

Schedule your job when:

- The job you are submitting is a “one-shot” job, but you don’t want it to run now. You want it to run at another, specific time, such as after a meeting or after hours.
- The job you are submitting should run regularly. This is especially useful for setting up a tape rotation scheme for your network.
- Your tape drive is currently busy and you want to run a backup job as soon as the drive is free. To do this, you would schedule your backup job with the current date and time. This way, when ARCserve discovers that the tape drive is busy, it will automatically retry the backup job until the drive is free.



Do the following to schedule your job.

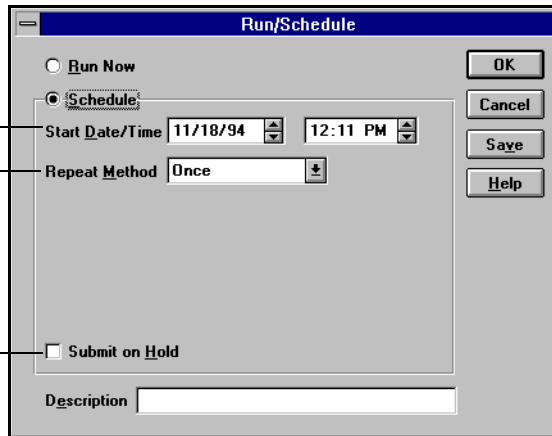
1. In the Run/Schedule dialog box, select Schedule.

The Run/Schedule dialog box appears, as shown below:

You must specify a date and time for the job to run. By default, the current date and time appear.


If you want the job to run regularly, specify the repeat method and interval.

Send the job to the queue with a "Hold" status. The job will not run until you use the Job Queue Manager to change its status to "Ready".



2. Specify the date and time you want the job to run .

The current date and time appear initially. Specify the time and date using one of the following methods:

- Type the specific date and/or time manually.
- Use the spin buttons  to select the date and time.
- Double-click in the date field to display a calendar to select a specific date.

Click to see the prior month.



Click to see the next month.

Highlight the date you want and click to select it.

3. (Optional) Specify a repeat method .

If you want this job to be run regularly, you need to provide the information in the *Repeat Method* fields.

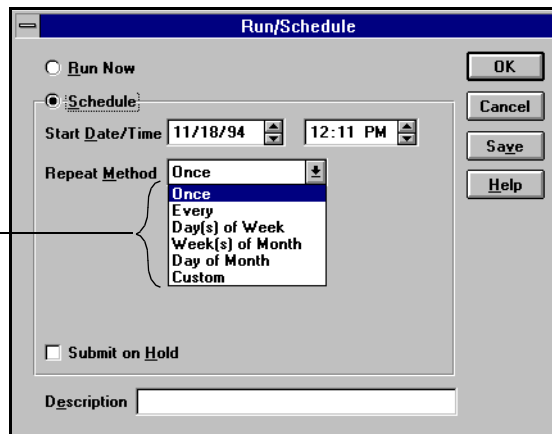
Specifying Repeat Method information

You can specify a repeat interval for a job once and then let ARCserve handle submitting the job to the queue for execution. This is especially useful when you are setting up a tape rotation scheme for your network. Set up the backup job, specifying your source machines, destination group and tape, backup method, and filters, and then specify when the job should run. This section describes how you use the Run/Schedule dialog box to specify a repeat interval for your job.

1. Specify the Repeat Method you want to use for the job .


By default, "Once" is selected. This means the job is scheduled to run only once. To change it, drop down the Repeat Method list, and select the method you want to use for the job.

Repeat methods available for a job. Depending on the method you choose here, fields may appear for you to specify more information.



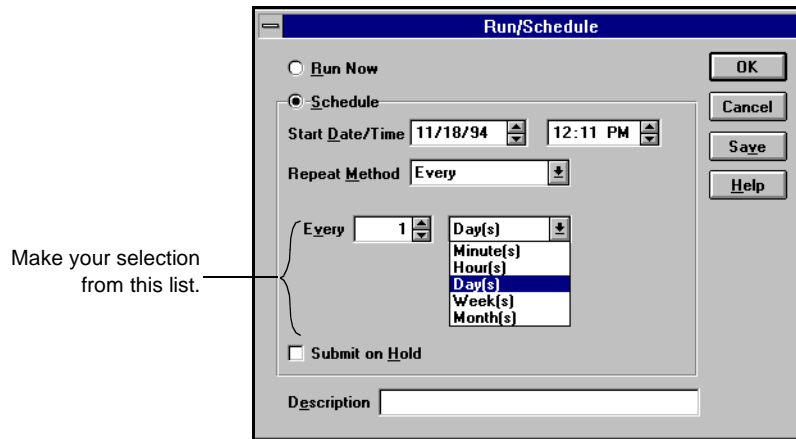
2. Specify the details required for the repeat method you selected.

Depending on the method you choose, extra fields are provided for you to specify the details for the repeat interval.



Every repeat method

If you select this repeat method, the Run/Schedule dialog box appears as such:



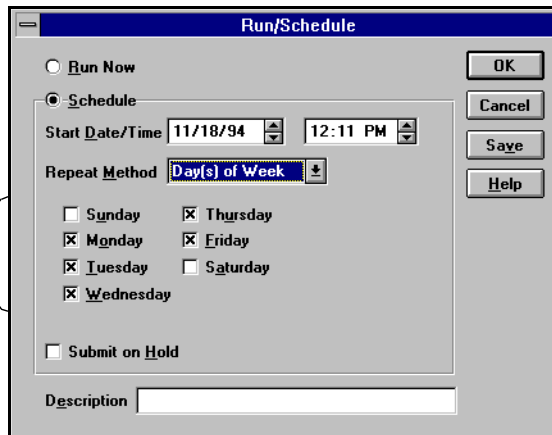
This repeat method should be set based on the following:

Every <i>n</i> minutes	The job will run every <i>n</i> minutes, beginning from the date and time you specified in the <i>Start Date/Time</i> fields. Specify the number of minutes.
Every <i>n</i> hours	The job will run every <i>n</i> hours, beginning from the date and time you specified in the <i>Start Date/Time</i> fields. Specify the number of hours.
Every <i>n</i> days	This is the default selection. The job will run every <i>n</i> days at the time you specified in the <i>Start Time</i> field, beginning on the date you specified in the <i>Start Date</i> field. Specify the number of days.
Every <i>n</i> weeks	The job will run every <i>n</i> weeks at the time you specified in the <i>Start Time</i> field, beginning on the date you specified in the <i>Start Date</i> field. Specify the number of weeks.
Every <i>n</i> months	The job will run every <i>n</i> months at the time you specified in the <i>Start Time</i> field, beginning on the date you specified in the <i>Start Date</i> field. Specify the number of months.

Day(s) of week repeat method

If you select this repeat method, the Run/Schedule dialog box appears as such:

Specify the day(s) of the week on which you want the job to repeat.


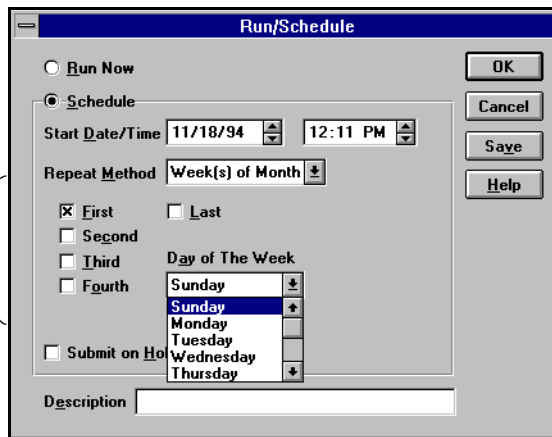


Use this repeat method if you want the job to run on specific day(s) of the week.

Week(s) of the month repeat method

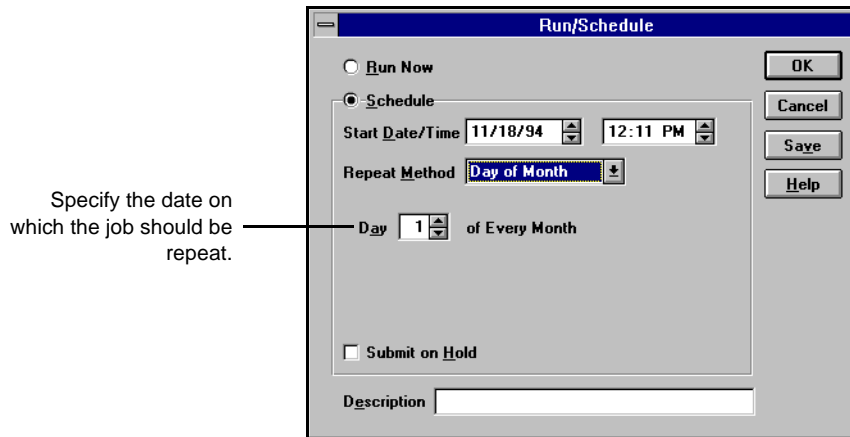
If you select this repeat method, the Run/Schedule dialog box appears as such:

Specify the week(s) of the month and the day on which you want the job to repeat.



Day of Month
repeat method

If you select this repeat method, the Run/Schedule dialog box appears as such:



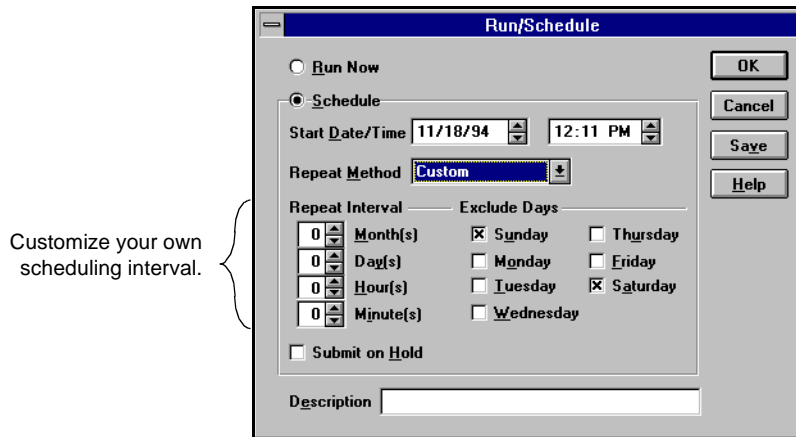
Use this repeat method for jobs that you want executed on the same date every month (say, the 15th of the month, or the last day of the month).



To have the job run every last day of each month, specify 31 in the *Day n of every Month* field.

Custom repeat method

If you select this repeat method, the Run/Schedule dialog box appears as such:



Use this method if you want the job to repeat at regular intervals (every day, for example) but you want to skip certain days (such as Saturday and Sunday).

The *Repeat Interval* should be set based on the following:

Every <i>n</i> months	The job will run every <i>n</i> months at the time you specified in the <i>Start Time</i> field, beginning on the date you specified in the <i>Start Date</i> field. Specify the number of months.
Every <i>n</i> days	The job will run every <i>n</i> days at the time you specified in the <i>Start Time</i> field, beginning on the date you specified in the <i>Start Date</i> field. Specify the number of days.
Every <i>n</i> hours	The job will run every <i>n</i> hours, beginning from the date and time you specified in the <i>Start Date/Time</i> fields. Specify the number of hours.
Every <i>n</i> minutes	The job will run every <i>n</i> minutes, beginning from the date and time you specified in the <i>Start Date/Time</i> fields. Specify the number of minutes.

Creating scripts for your jobs

Save your job settings like you'd save a file

You can save any type of job as a script. The script will contain the source, destination, and options that you selected, as well as the schedule information. 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 NT machine running ARCserve for Windows NT
- You can quickly re-submit regularly executed jobs after a job has been accidentally deleted

To use a script, go to the Job Queue Manager and use the Add Job button. When you specify the job script name, all of the information that you set when you saved the job script will be used, including the files and directories and the scheduling information.

Creating a job script

Follow these directions to create a script:

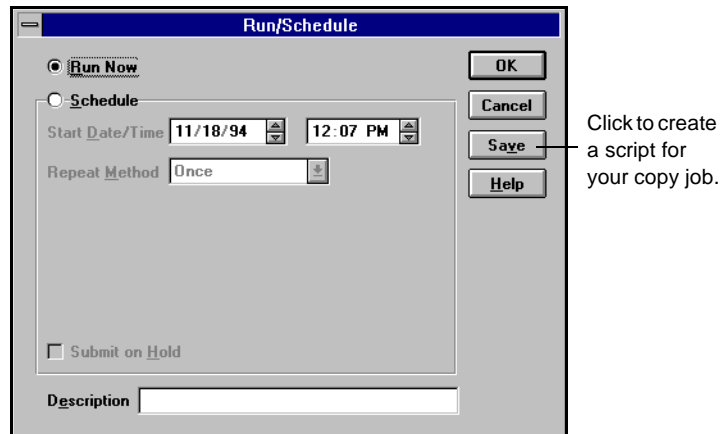
1. Set up the job.

Select the source, destination (if required), and any options or filters. Refer to the appropriate chapter in this manual for more information about setting up a job.

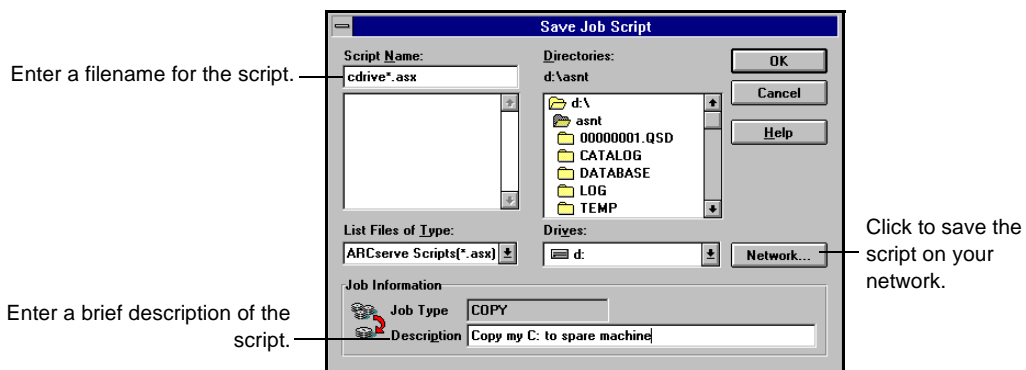
2. Click the Run/Schedule button.



The Run/Schedule dialog box appears, as shown below:



3. Specify scheduling information for the job.
Refer to previous sections in this chapter for details about specifying scheduling information.
4. Click the Save button.
You can also save a script by selecting *Save As* or *Save* from the File menu.
The Save Job Script dialog box appears.



5. Give the script a name.

By default, ARCserve scripts can be stored in the ARCserve home directory with a default extension “.asx”. You can, however, specify a different directory in which to save your script.

6. (Optional) Enter a description of the job .

This can help you identify the correct script at a later date.

7. Click OK.

To use a script -
add it to the
queue via the Job
Queue Manager

Follow these directions to use a script:

1. Go to the Job Queue Manager.
2. Click the Add Job button.
3. Specify the script file name.

Refer to Chapter 8, “Managing Jobs and the Job Queue”, for more details about using the Job Queue Manager and the Add Job button.

8

Chapter

MANAGING JOBS AND THE JOB QUEUE

Information about all jobs is stored in ARCserve's job queue.

In this chapter, you will learn:

Page

- | | | |
|------|---|--|
| 8-2 | ➤ | About the ARCserve job queue |
| 8-3 | ➤ | How to use the Job Queue Manager |
| 8-6 | ➤ | How to manage jobs in the queue |
| 8-13 | ➤ | The difference between stopping and deleting an active job |
| 8-17 | ➤ | How to monitor an active job |

Job queue basics

How does
ARCserve's job
queue work?

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. (The scanning frequency can be changed through the ARCserve Server Admin. Refer to Chapter 14 for more information.) At the appropriate date and time, the job is executed.

The Job Queue
Manager window


You can do the following from the Job Queue Manager:

- Add a job
- Change job execution date, time, or status
- View contents of a job
- Modify a job
- Delete a job
- Monitor an active job
- Stop an active job

The remainder of this chapter explains how to use the Job Queue Manager.

Opening the Job Queue Manager window

All job queue operations start from the Job Queue Manager. To open the Job Queue Manager:



1. From the Quick Access window, click Job Queue.

The Job Queue Manager displays a list of all jobs in the queue. Information about each job appears on its own line.

Status	Execution Time	Job Type	Last Result
DDONE	<Run Now>	Backup	Completed
63%	<Run Now>	Backup	Completed
20%	<Run Now>	Count	Completed
ACTIVE	<Run Now>	Backup	Failed
READY	6/20/95 11:35 PM	Count	
READY	6/21/95 10:28 PM	Backup	
HOLD	6/22/95 6:30 PM	Restore	

Summary Panel:

- Job Type: Backup
- Status: DDONE
- Last Result: Completed
- Owner: MichaelT
- Execution Time: <Run Now>
- Submitted Time: 6/20/95 10:19 PM
- Description: Document backup
- Source: \\MIKET351\Z\$
- Destination: MIKET_TAPE IN GROUP1

When you drag this bar you will see additional fields of information:

Job Type - Backup, Restore, Copy, etc.

Last Result - The result of the last job execution. Refer to Page 8-7 for more information.

Owner - Name of the person who submitted the job.

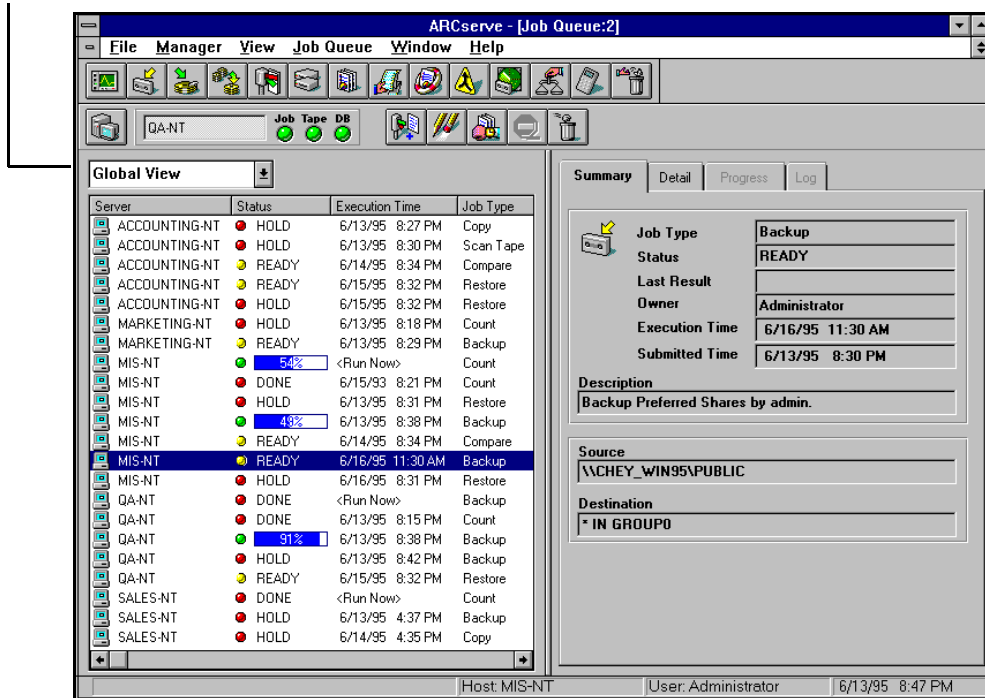
Description - A user-defined description of this job.

Changing to Global View


By default, the Job Queue window displays the status of jobs on the selected server. To view the status of all jobs from multiple ARCserve Servers, change to Global View. Global View displays the jobs from the ARCserve Servers that are in the Preferred Servers list for the Job Queue window.

To change to Global View, select *Global View* from the drop-down list box above the job status list, as shown:

Select Global View




Using the hidden menus If you highlight a job listed in the queue and click your right mouse button, the following menu will appear:



Each of these functions is discussed in this chapter.

If you are viewing information in the Detail or Log tab and you click your right mouse button, the following menu will appear:




Managing jobs in the queue

You can manage jobs that are in the queue with the Job Queue Manager. Job management includes the following tasks:

- Reviewing information about the last time jobs that have been executed
- Adding a previously saved script to the queue
- Changing a job's execution date, time, or status (this includes resubmitting a job)
- Displaying information about a job
- Modifying information about a pending job
- Deleting a job
- Monitoring an active job
- Stopping an active job

Each function listed above is explained on the following pages.



Reviewing information about executed jobs

A job listed in the Job Queue with a *Done* status is a job with no repeat interval that has already been executed.

Done 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 Chapter 14 for more information.

A repeating job (a backup job that runs every Friday, for example) will not have a *Done* status after it has been executed. Instead, it will have a *Ready* status (it is ready to be executed again next Friday).

The *Last Result* field

The *Last Result* field tells you whether or not your executed job was successful. And, if it wasn't successful, the *Last Result* field helps you figure out why the job may have failed.

Last Result	Explanation	Possible Cause s
Completed	All of the nodes and drives/ shares were processed.	n/a
Incomplete	Some nodes and/or drives/ shares were missed.	A machine was turned off, a share was not available, or some files were missed.
Cancelled	The user cancelled the job.	Someone cancelled the job from the Job Queue, answered NO or CANCEL to a prompt, or the job required the user to answer OK or to insert a tape but the timeout was reached and the job was cancelled. (Timeout is set in the Tape options in the Backup Manager window.)
Failed	The job was started but could not be completed.	There was an error reading from or writing to the tape, the tape group was already open, there was no tape in the drive, or a memory allocation error occurred.
Run Failed	The job was started, but the program that runs the job (the Job Runner) failed.	There was not enough memory to run the job or a DLL file was not found.

Last Result	Explanation	Possible Cause s
Crashed	The job was started, but before it was completed, something happened to the Job Runner and Job Runner could not update the job queue.	A system error occurred in the job, such as a memory violation.

What can I do if my job didn't complete successfully?

If you determine that your job was not successful, you can use the information provided in the *Last Status* field to remedy the problem. You can then resubmit the job to run. (Refer to Page 8-11 for information about resubmitting a job.)

Before going home last night, Bryan scheduled a backup of his machine. When he checked the Job Queue Manager this morning, the Last Status field said Cancelled. Bryan quickly turned to the "Managing Jobs and the Job Queue" chapter to find out what could have caused his job to fail. After reading the possible causes for a Cancelled Last Result, Bryan did some investigative work and realized that the tape he used was nearly full and ARCserve was looking for a second tape. Since he was not here to put a new tape in, the timeout had been reached and the job couldn't continue. Bryan found a new tape and resubmitted his job.

The Queue Log

The Queue Log displays information about jobs that have been executed. The level of detail in this log is determined when you set up a job. You can select 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). You also have the ability to disable reporting.

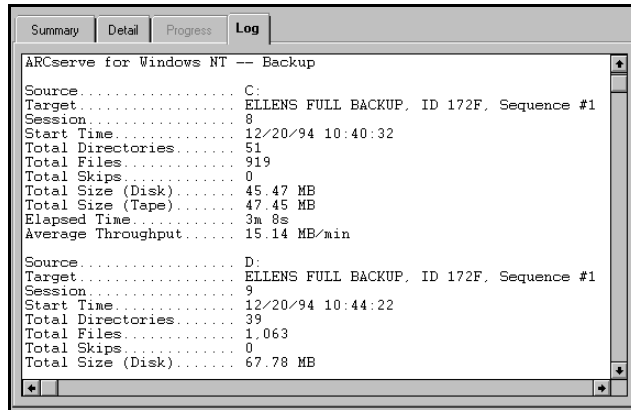
To display the Queue Log:

1. Highlight the job you want.

2. Select the Log tab.

The log is displayed:

This is a sample Queue Log that displays summary information only.



Adding a job to the job queue

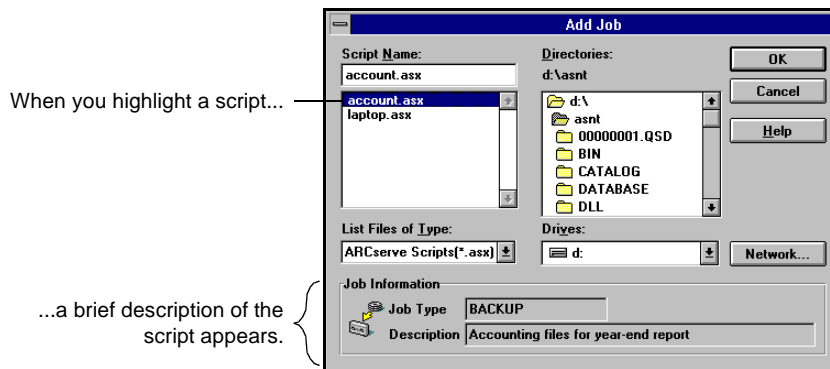
You can quickly submit a job to the queue by using a previously saved script. A script is a job that you saved to a file. It contains the original source, destination, option, and schedule information for the job. Refer to the 'Creating scripts for your jobs' section on page 7-28 for more information about scripts.

Follow the steps below to add a job to the job queue using the Job Queue Manager.

1. From the Job Queue Manager, click the Add Job button.



The Add Job dialog appears.




2. Select the file name for the script you want to submit to the queue.

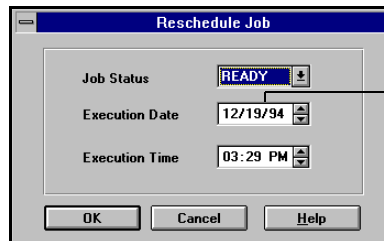
The script that you select is added to the job queue.

Changing a job's execution date, time, or status

The Job Queue Manager makes it easy for you to 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.

Follow the directions below to change execution information about a job:

1. Highlight the job you want to change.
You can only change a job's execution date, time, or status if the job is not active.
2. Click the Reschedule Job button. 
3. Change the job status, execution date, and/or execution time.



You can double-click in the date field to display a calendar to help you select a specific date...

...Highlight the date you want and click to select it.



Job status can be changed to "Hold" or "Ready". "Ready" job status indicates the job will execute at the specified date and time. A job with "Hold" status will not execute until you change the status to "Ready".

After clicking OK, the job is changed to the status, execution date, and execution time you specified.

Displaying information about a job

The Job Summary log

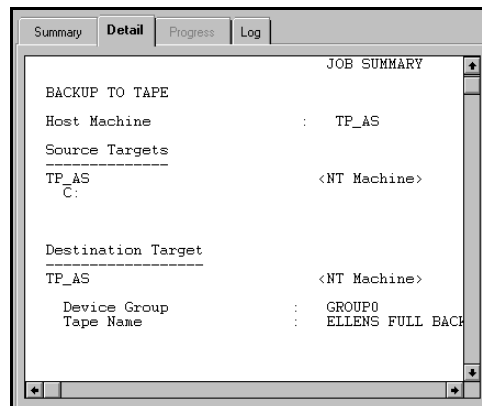
The Job Summary log displays the details about a job, including the source and destination target(s) and tape statistics, such as sequence and session number.

To display the Job Summary log for a job:

1. Highlight the job you want to view .
2. Select the Detail tab .



The log is displayed.

This is a sample section from a Job Summary log.




Stopping an active job

You can use the Stop Job button or the Delete Job button to cancel an active job. However, if the job you are stopping is a job that repeats at regular intervals (this is determined when you create the job), you should read the following information to determine which button to use:

- Stop Job button  - cancels the active job and reschedules it for its next regular interval.
- Delete Job button  - cancels the active job and deletes it completely. The job will not be rescheduled for its next regular interval.

To stop an active job:

1. Highlight the active job you want to cancel.
2. Click the Stop Job button. 
3. Click OK to confirm.

Modifying a job

Use the Job Queue Manager to modify a job already in the queue. This way you can add options or additional sources to an existing job, without having to create a new job.

Follow the steps below to modify a job:

1. Highlight the job you want to change.

You cannot modify an active job. Only jobs with a Ready, Hold, or Done status can be modified.

2. Click the Modify button. 

The job that you originally created appears in its appropriate window, with the original source(s), option(s), and destination(s) selected. For example, if the job is a backup job, the Backup Manager appears. If the job is a restore job, the Restore Manager appears.

3. Make the desired changes.
4. Click the Run/Schedule button.
Change or re-confirm the original execution information.
5. Click OK.


Deleting a Job

You can delete jobs from the queue, whether they are active, done, or pending. Deleting jobs from the Job Queue Manager window will completely remove them from the job queue. This means that repeating jobs will not be rescheduled.



You may not want to use the Delete Job button if you are deleting an active job that repeats at automatic intervals (determined when you create the job). The Delete Job button interrupts and removes the job completely from the queue. As a result, the job is not rescheduled for its next automatic interval. You will have to recreate the job if you did not save it to a script file.

To delete a job:

1. Highlight the job you want to delete.
2. Click the Delete Job button. 

After confirming, the job is removed from the Job Queue list.

Changing the way jobs are listed in the job queue

When you first open the Job Queue Manager window, the jobs in the queue are listed in order of execution time. At any time, you can change the way in which jobs are listed in the queue. Jobs can be sorted by clicking any of the following fields:

- Status
- Execution Time
- Job Type
- Server (for Global View)

Sorting the job queue is for informational purposes only. It does not affect the order in which jobs are processed.

Follow the step below to sort the job queue:

1. Click on the title of the field you want to sort.
Jobs are sorted alphabetically or chronologically within the field you selected.

Monitoring an active job

You can monitor the real-time status of a job from the:

- Job status box on left side of window (refer to 'Opening the Job Queue Manager window' section on page 8-3)
- Progress window (described below)

To monitor an active job from the Progress window:

1. Highlight the active job you want to monitor.
2. Select the Progress tab.

The Progress window is displayed. Depending upon the type of job that is running, some of the fields may not be used. For example, the *Remaining Time* field will not be visible for all job types except for Backup.

Information about the session ARCserve is currently processing.

Tape name, ID, tape sequence and session number.

Summary information for the session ARCserve is currently processing and the estimated time remaining.

Percentage completed.

Total Files	MB Processed	Elapsed Time
419	21.76	1m 18s
MB/Minute	MB Estimated	Remaining Time
16.74	45.47	1m 27s



9

Chapter

MANAGING TAPES & TAPE DEVICES

This chapter contains information about tapes and tape devices.

In this chapter, you will learn how to:

<i>Page</i>	
9-6	➤ Get information about tape devices and the tapes that are in them
9-12	➤ Configure tape groups
9-18	➤ Perform various maintenance functions on tapes, such as formatting, erasing, and retensioning

Device Manager basics

The Device Manager is informational as well as functional

Whenever you want information about tape devices that are connected to your system, the tapes in these devices, or the status of these devices, turn to the Device Manager.

Whenever you need to format, erase, or retention tapes, use the Device Manager. Also use the Device Manager to change the compression mode for your tape drives, if they support compression.

Several tape devices = tape groups

If you have more than one tape device connected to your Windows NT machine, ARCserve lets you separate them into two or more “groups”. Establishing tape groups is a key component to ARCserve’s flexibility and efficiency.

Tape groups allow for:

- Parallel streaming - you can have several operations occurring simultaneously, one at each tape group configured for your system.
- Automated tape spanning - if, during a backup session, the end of a tape is reached, ARCserve automatically spans to another tape in another device within the same tape group. Note that this requires two or more devices in the tape group for this to happen.

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

Multiple tape groups = parallel streaming

If you have more than one tape group configured on your machine, you can take advantage of “parallel streaming”. Parallel streaming means that you can have more than one job requiring a tape device running at a time. For example, suppose you have two tape groups, “GROUP1” and “GROUP2”, respectively. You could then have a backup job running with GROUP1, while a restore job runs with GROUP2.

Multiple devices in a group = “automated” tape spanning


Tape spanning means that, when one tape becomes full, the session will automatically span to another tape. If you have a tape group that consists of two or more devices, tape spanning will occur across devices. That means, when one tape in a group drive becomes full, the job spans to another tape in another device within the same tape group.

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

How do I decide which devices can be grouped together?

There is only one rule when assigning tape devices to the same group: the devices in the group *must* be of the same make, model.

Tape groups allow you to take advantage of parallel streaming. Having several devices in one group allows you to take advantage of "automated" tape spanning.



Refer to 'Configuring tape groups' section on page 9-12 for more information about setting up and configuring tape groups.

The rest of this chapter describes how to get information about tape devices, how to interpret this information, and how to maintain your tapes.

Using the Device Manager

The starting point for all tape and tape device operations

The Device Manager is the starting point for all tape and device monitoring and maintenance operations. The Device Manager is shown below. Each device option is discussed in more detail throughout the rest of this chapter.

To access the Device Manager:



1. Click the Device Manager button from the Quick Access window.

The Device Manager appears:

Mini Server Admin.

Configure device groups.


Buttons to format, erase, retention, compress data, and eject tapes.

Tree view of adapters, devices, and tapes is displayed here.

Summary and detailed information about what you selected in the left pane.

Adapter Information	
Adaptec Inc.	
Description: Aha154x	
Board Number	0
I/O Port	0330
SCSI ID	7
Memory Address	N/A
Bus Type	ISA
DMA	5
IRQ	11

Host: TP_AS User: Administrator 12/21/94 2:20 PM



Getting information


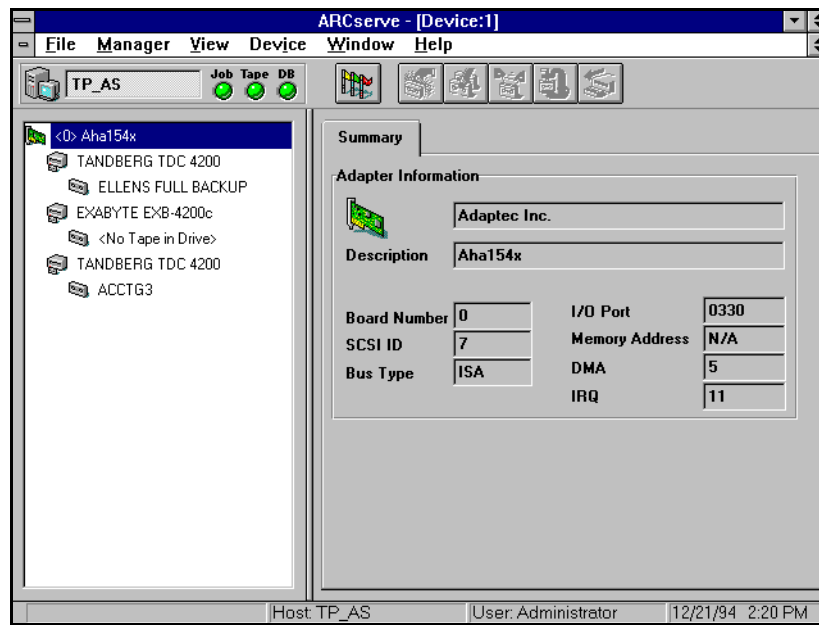
You can view a variety of information about the tapes you use and the tape drives connected to your system using the Device Manager. The rest of this section describes:

- Getting general or detailed information about an adapter card.
- Getting general or detailed information about a tape device.
- Getting general or usage information about a tape in a tape device.

Getting information about an adapter card

To get information about an adapter card, follow the steps below:


1. Highlight the adapter card.



Getting information about a tape device

When you highlight a tape device, you can see two types of information: summary and detail. To view either type of information for a tape device, follow the steps below:

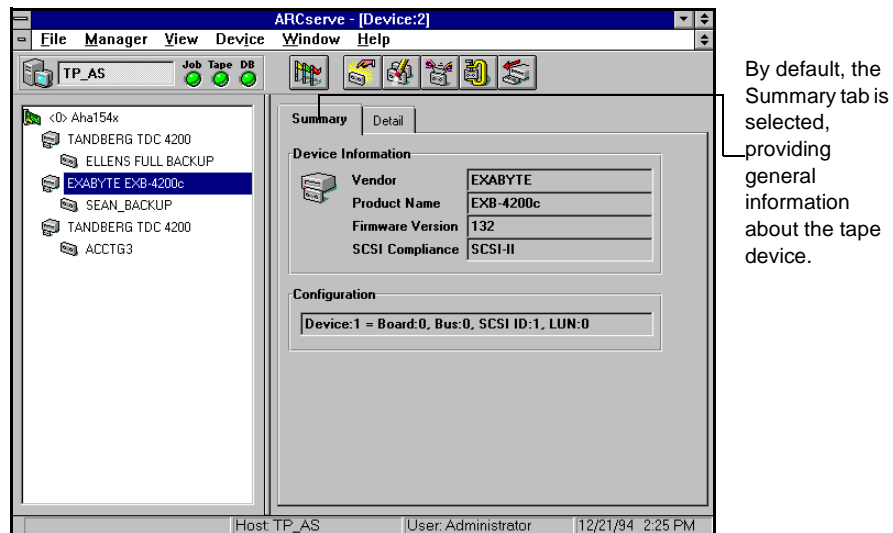
1. Highlight the tape device you want information about.



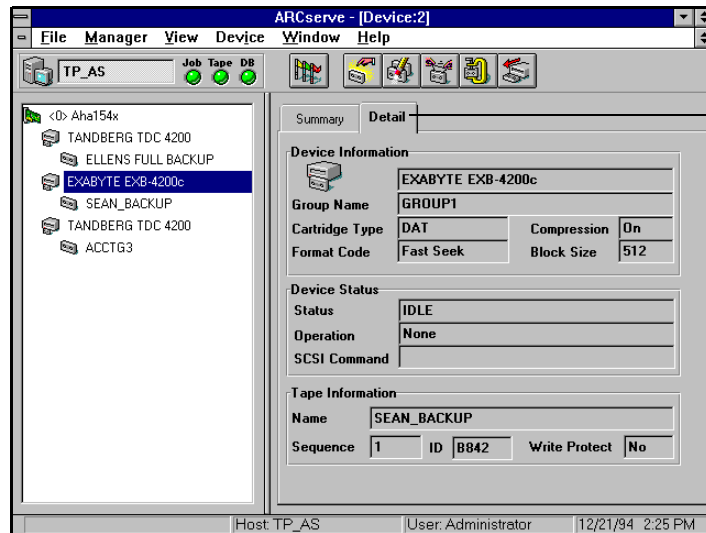
2. Select the Summary or Detail tab.

Summary information

Selecting the Summary tab displays general information about the tape device, such as its vendor, product name, firmware version, and SCSI compliance.



Detailed information Selecting the Detail tab displays detailed information about the tape device.




Selecting the Detail tab shows specific information about the tape device you selected.



Getting information about a tape in a tape device

When you highlight a tape, you can see two types of information: summary and details. To view either type of information for a tape, follow the steps below:

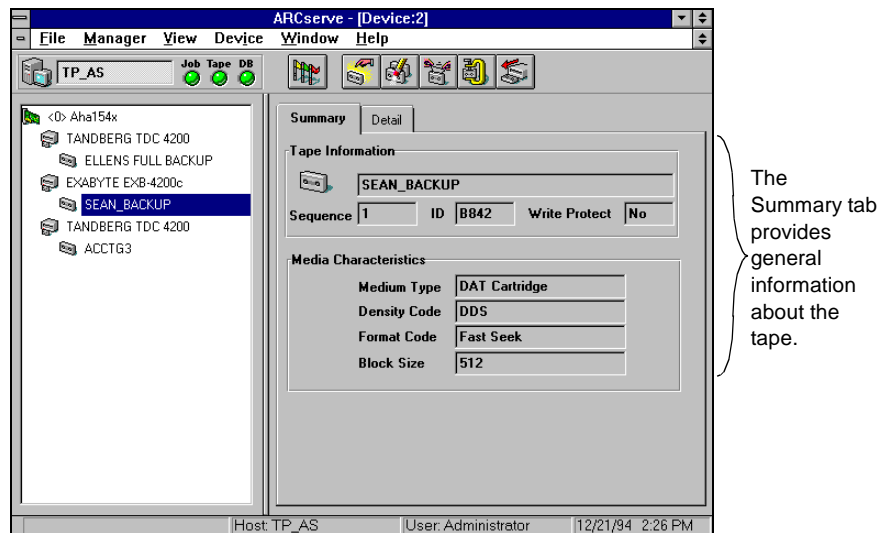
1. Highlight the tape you want information about.



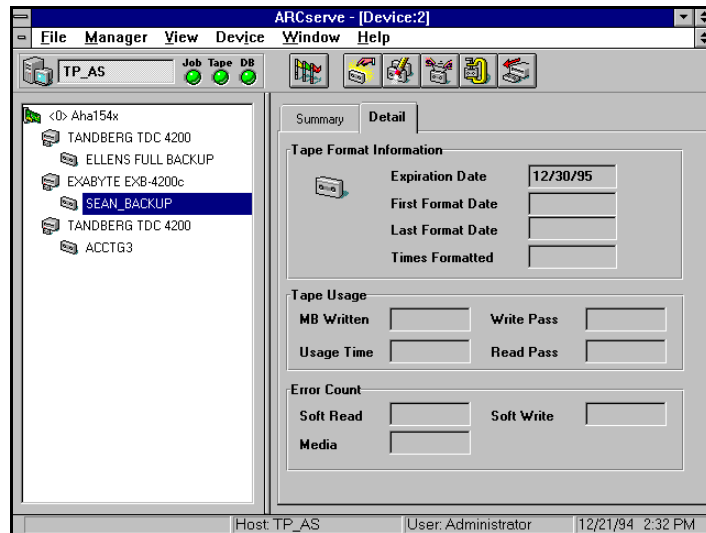
2. Select the Summary or Detail tab.

Summary information


Selecting the Summary tab displays general information about the tape, such as the tape name, sequence number, ID, and whether it is write protected. Media characteristics are also displayed.



Detailed tape information Selecting the Detail tab to see detailed information about the tape's usage.



The Detail tab provides specific information about the tape, such as format, usage, and read/write error information.



Configuring tape groups

ARCserve uses “tape groups”

As discussed at the beginning of this chapter, ARCserve provides the concept of “tape groups”. ARCserve allows you to separate your tape devices into groups. Grouping devices allows you to take advantage of parallel streaming. And, if you have several devices in a group, you can also take advantage of “automated” tape spanning.

One rule for grouping multiple devices in a group

If you want to have several devices in a group, there is only one rule. Each device in the group must be the same make and model.

When ARCserve’s Tape Engine is started, each tape device that you have is automatically placed in a tape group. Identical drives are grouped together and differing 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 tape devices, you can keep these tape group names, or, you can regroup and rename them. After you start ARCserve, you can use the Device Manager and the instructions in the following sections to:

- Create a new tape group
- Assign a device to a tape group
- Remove a tape device from a tape group
- Rename a tape group

Creating a new tape group

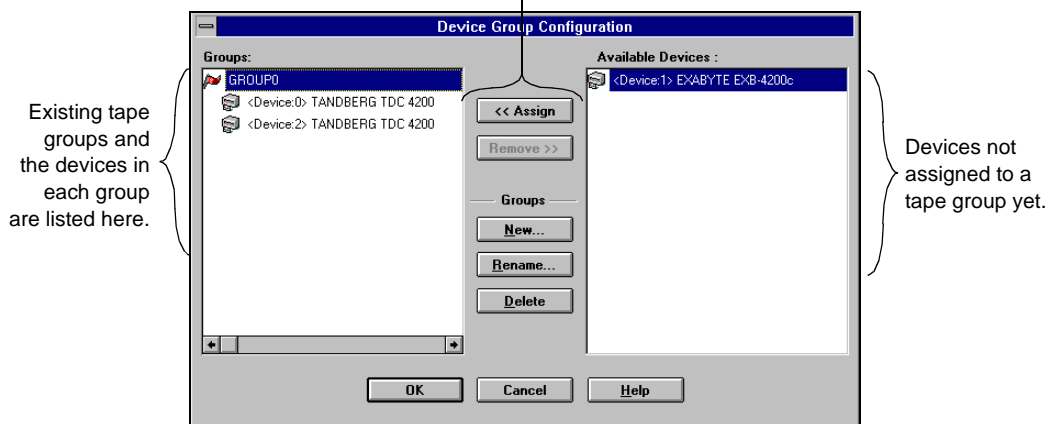
Follow these directions to create new tape groups and assign tape devices to them:

1. Click the Device Group Configuration button .



The Group Configuration dialog box appears:

Use these buttons to assign and remove devices from groups, create a new tape group, rename an existing tape group, or delete a tape group.



2. Click the New... button.

The New Group dialog box appears.

3. Enter a name for the tape group and click OK.

The new tape group appears in the Groups box. Now you can begin assigning tape devices to this tape group.

Assigning a tape device to a group

Follow these directions to add tape devices to a tape group:

1. In the Device Group Configuration dialog box, select the device from the Available Devices box.

The Available Devices box lists all the tape devices that are not assigned to a tape group. From this list, select the one you are assigning to the group.

2. From the Groups box, select the group you want the device assigned to.

Highlight the name of the tape group.

3. Click the <<Assign button.

The device is removed from the Available Devices box and is placed beneath the tape group in the Groups box.

4. Repeat steps 2 and 3 to assign more devices to groups.
5. When finished, click OK.

Removing a tape device from a tape group

Follow the steps below to remove a tape drive from a tape group:

1. In the Device Configuration dialog box, highlight the device to remove.

Devices are listed below their assigned tape groups in the Groups box. Select the one you want to remove by highlighting it.

2. Click the Remove>> button.

The tape device is removed from under the tape group in the Groups box and appears in the Available Devices box.

3. Repeat steps 1 and 2 to remove more devices from tape groups.
4. When finished removing tape devices, click OK.

Other tape group configuration features

Deleting a tape group

Follow the steps below to delete a tape group.

1. Remove the tape devices assigned to the tape group .
Refer to 'Removing a tape device from a tape group' section on page 9-16 for details on performing this step.
2. Select the tape group you want to delete.
Highlight the name of the tape group in the Groups box.
3. Click the Delete button.
The tape group is removed from the Groups box.

Renaming a tape group

If you want only to change the name of a tape group, follow the steps below:

1. Select the tape group you want to rename.
Highlight the name of the tape group in the Groups box.
2. Click the Rename button.
The Rename Group dialog box appears.
3. Enter a new name for the tape group and click OK .
The new tape group name is reflected in the Groups box.

Maintaining your tapes

Tape
maintenance
tasks you can
perform

You can perform the following maintenance tasks on your tapes:

- Formatting
- Erasing data
- Retensioning
- Compressing data - this task can only be used if the tape drive supports tape compression
- Ejecting from the tape drive - this task can only be used if the tape drive supports this feature



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

The rest of this section shows you how to use the options listed above.

Formatting a tape

Format your tapes before you use them


Although ARCserve automatically formats blank tapes, you can use this option to manually format your tapes. Formatting writes a new label at the beginning of the tape, effectively destroying all existing data on the tape.

Low level formatting, as is required on most hard drives and some mini cartridge tape drives, is not required for drives that ARCserve for Windows NT supports.




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

To format a tape:


1. Insert the tape into a tape drive.
2. In the Browser, highlight the tape.
3. Click the Format button. 

The Format dialog box appears.

4. Enter a name for this tape and select a date for this tape to be put out of service.



A tape's life is generally based on "passes". A pass is defined as the tape drive head "passing" over a given point



on the tape. For example, a backup without verification constitutes one pass, whereas a backup with verification constitutes two passes.

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

You should choose an expiration date based on how you will use the tape. If you plan to use the tape 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 tape only once or twice a month, you can set the expiration date to two or three years from the current date.

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

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



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

5. Click OK to start formatting the tape.

A message appears asking you to confirm formatting the tape.

Erasing a tape

Use this option to erase all data from a tape. ARCserve also erases all references to the contents of this tape (if any) from the database. When and if you re-format this tape, its physical history (read and write passes, etc.) will be carried over.



Use this feature with care! Once you erase data on a tape, the data is gone.

9

There are two erase options from which you can choose: Quick Erase and Long Erase.

Quick Erase

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


Long Erase

A Long Erase completely removes all data from a tape. It takes much longer than a Quick erase, but the tape is considered blank, as if it were just formatted. For security reasons, if you want to make sure that the data on a tape is gone completely, use Long Erase.

To erase a tape:

1. Insert the tape that you want to erase into a tape device.
2. Select the tape.

Refer to Chapter 3, "Basic Procedures" for details on selecting a tape.

-
3. Click the Erase button. 

The Erase dialog box appears.



4. Select Quick Erase or Long Erase.
5. Click OK to start erasing data from the tape.

Retensioning a tape

Computer tapes are similar to audio tapes in that the more you use them, rewinding and fast forwarding, the more the tape becomes uneven and loose on the spools. When a tape becomes unevenly wound, it is prone to errors, may jam, or worse yet, break.

Use ARCserve's Retension feature to make sure a tape is evenly wound and properly tensioned. Retension a tape especially if you are having trouble writing to it or reading from it.




This feature applies only to Quarter Inch Cartridge tapes. It does not apply to 8mm or 4mm Digital Audio Tapes (DATs).

To retension a tape:

1. Insert the tape into a tape device.
2. Select that tape.

In the left pane of the Device Manager, expand the tree under the device that the tape is in, and highlight the tape.

3. Click the Retension button. 
4. Click OK to start retensioning the tape.

Compressing data on tapes


If your tape drive supports compression, you can tell ARCserve to use it by turning it on.

Normally, you will want to leave compression on. The only time you should turn it off is if you plan to use the tape in another drive that doesn't support compression. The drive that doesn't support compression will not be able to read the compressed data on the tape.



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

To turn compression on or off:


1. Select the tape device that supports data compression.
In the left pane of the Device Manager, highlight the tape device.
Select the Detail tab in the right pane to check that compression is supported by this device.
2. Click the Change Compression Mode button. 
3. Click OK to set the compression mode to Off (if it is On) or On (if it is Off).

Ejecting tapes

You can eject tapes from a drive using ARCserve's Device Manager, if your device supports this feature.

To eject a tape with ARCserve:

1. Select the tape device you want to eject the tape from .
In the left pane, highlight the tape device.

2. Click the Eject button. 

The tape is ejected from the tape drive.



10

Chapter

MANAGING THE ARCSERVE DATABASE

This chapter explains how ARCserve uses its database to store information about the jobs you run through ARCserve, tapes you use with ARCserve, and the devices you use for backups.

In this chapter, you will learn:

<i>Page</i>	
10-2	➤ About the Database Manager
10-5	➤ About the Summary View
10-6	➤ About the Job View
10-9	➤ About the Tape View
10-12	➤ About the Device View

Database basics

ARCserve stores information about the jobs you have run, the tapes you have used for backups, and the devices you use for backups.

How is database information used in ARCserve?

Database information helps when you restore files from tape. The ARCserve database stores information about every file and directory that has been backed up. When you are looking to restore a specific file, ARCserve's database helps you determine which tape to use. This eliminates the time it would take to search through tapes looking for files.

ARCserve's Reports Manager is another place where database information is used. Information for the reports comes from the ARCserve database. The Reports Manager window allows you to view and print reports with this information. Refer to Chapter 12 for more information.

What can you do in the Database Manager window?

You can do the following in the Database Manager:

- Keep track of the location of your tapes
- Determine the session number of your most recent backup
- Determine if a tape should be retired
- View log information about jobs you have run
- Delete old records from the database

The Database Manager window also displays graphical information about the hard disk on your machine and the amount of space the database takes up.

The remainder of this chapter describes each database view and explains how to use the Database Manager window.

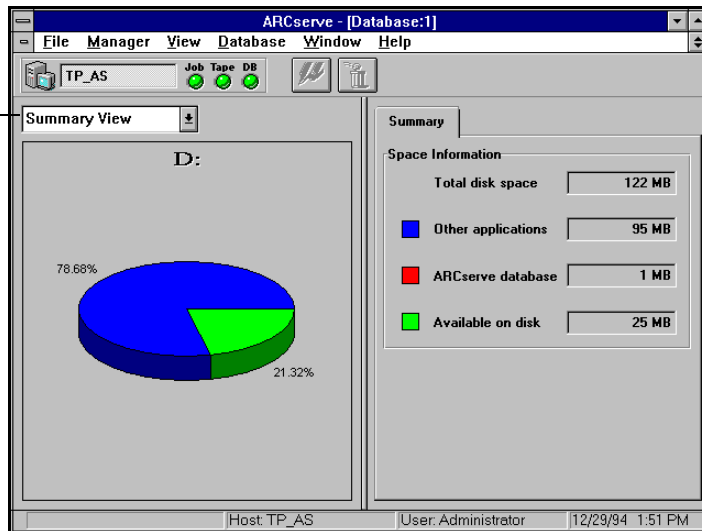
Opening the Database Manager window

All database operations start from the Database Manager window. To open the window:



1. Click the Database button from the Quick Access window .
The Database Manager Window is displayed. By default, you are looking at summary information for your hard disk (where ARCserve is installed).

Select another database view from this menu.



10

Selecting a database view

There are four database views from which you can choose:

- **Summary View** - A summary of the percentage of hard disk space being used by the database, all other files, and how much space is free (where ARCserve is installed).
- **Job View** - A list of all jobs that you have run through ARCserve including details about sessions and individual files.

-
- **Tape View** - A list of all tapes that you have used with ARCserve including tape statistics and session information for each tape.
 - **Device View** - A list of all devices you have used with ARCserve including SCSI ID, device statistics, and errors.

The next few pages describe each database view in detail.

Sorting records in a database

When you first open a Database window, the records appear in the following order:

- Job ID (Job View)
- Tape ID (Tape View)
- Adapter number (Device View)

You can sort the records by any field in the database.

To sort the records:

1. Click on the field name on which you want to sort.

Deleting a record from the database

In addition to setting ARCserve to automatically prune old database records (refer to Chapter 14, “Administering the ARCserve Server” for more information), you can manually delete database records. You may want to do this if you’ve replaced a tape device, for example. To delete a record:

1. Highlight the record(s) you want to delete.
You cannot delete the last Job View record.

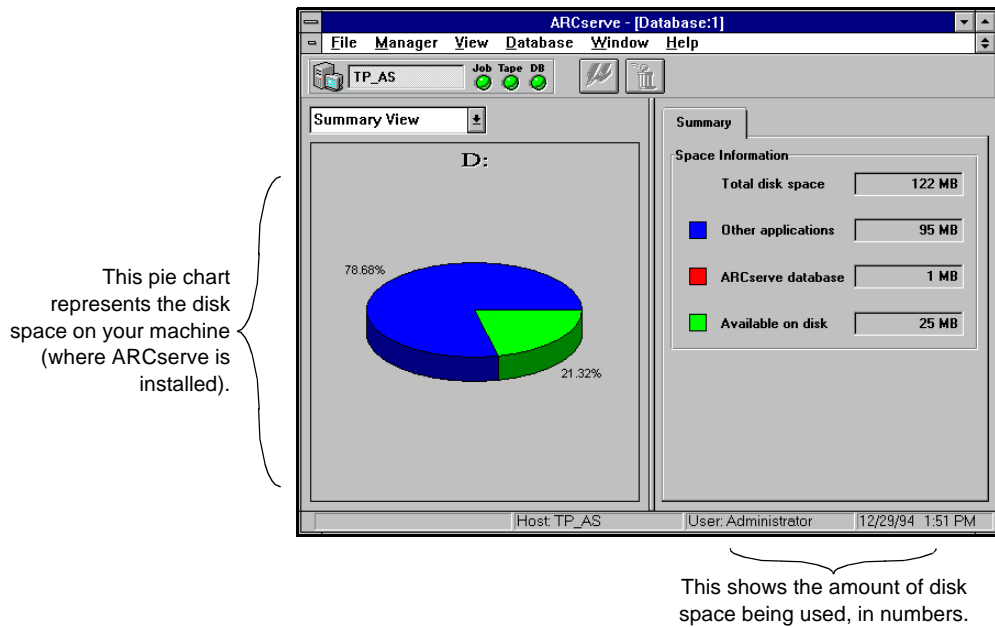
2. Click the Delete button. 

3. Confirm the deletion.

Make sure this is the record you want to delete.

Summary View

The Summary View is displayed when you first open the Database Manager window. This screen is informational only. You cannot modify it.



10

This window displays information about your machine's hard disk and the size of your ARCserve database.



If your database has grown very large, you may want to set ARCserve to prune old records from the database. Refer to Chapter 14, "Adminstrating the ARCserve Server" for more information.

Job View

The Job View window displays information about all jobs that have been processed by ARCserve.

Each time ARCserve runs a job, a record is added to the database. This record contains basic information about the job, such as type of job, who ran it, when it started, when it ended, and its status (finished, incomplete, etc).

For backup jobs, session and file information are all recorded in the database.

The screenshot shows the ARCserve Job View window. The main pane displays a table of jobs:

ID	Type	Status	Start Time
1	Backup	Finished	12/20/94 10:40:24
2	Backup	Finished	12/20/94 12:37:56
3	Backup	Failed	12/20/94 12:36:38
4	Backup	Finished	12/20/94 15:21:22
5	Count	Finished	12/20/94 15:20:48
6	Count	Finished	12/20/94 15:37:22
7	Backup	Finished	12/20/94 15:42:50

Below the jobs table is a 'Sessions' table:

Ses.#	Type	Path	Status
8	NT FAT	\\TP_AS\C:	Finished
9	NT NTFS	\\TP_AS\D:	Finished
10	ARCserve Dat.	\\TP_AS\D:\ASNT\D:	Finished

The right pane shows 'Job Information' for the selected job (ID 1):

Job Information
ID: 1
Type: Backup
Status: Finished
Start Time: 12/20/94 10:40:24
End Time: 12/20/94 10:52:52
Submitter: Administrator
Description:
Tape
Name: ELLENS FULL BACKUP
ID: 172F Sequence No. 1

Annotations:

- Highlight a backup job to display session information below and to display job information in the right pane.
- If you highlight a session, you can select the Session or Files tab to see additional information about that session.

Viewing details about a job



You can view log information for each job. In addition, for each backup job, you can see a list of the sessions that were created on tape and a list of files backed up for each session.

To view log information:

1. Highlight a job.

2. Click the Log tab.

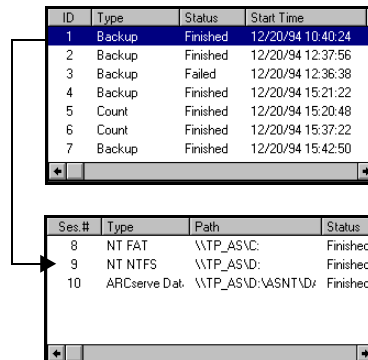
The log displays information about the job. The level of detail is determined when you first set up the job.



To view session information for a backup job:

1. Highlight a backup job.

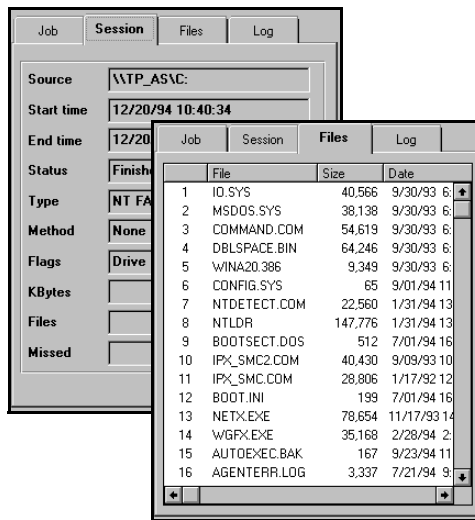
Session information is displayed in the *Sessions* section.



2. Highlight the session you want to know more about.

-
3. Select the Session tab or the Files tab to see additional information about the session.

Select the Session tab to see information about the backup session or select the Files tab to see a listing of files that were backed up.



Tape View

The Tape View window displays information about the tapes used with ARCserve. This includes information about formatting, how much the tape has been used since it was put into service, and number of errors (if any) that have occurred while ARCserve was using the tape.

Each time a tape is formatted by ARCserve, a new record is added to the database.



Highlight a tape to display session information and statistics below and to display tape information in the right pane.

If you highlight a session, you can select the Statistics, Session, or Files tab to see additional information about that session.

ID	Tape Name	Seq #	Last Format Date
1	ELLENS FULL BACKUP [C	1	11/15/94 15:16
2	ACCTG3 [ID:F5A0]	1	11/03/94 9:37

Ses.#	Type	Path	Status
8	NT FAT	\\TP_AS\C:	Finish
9	NT NTFS	\\TP_AS\D:	Finish
10	ARCserve Dat.	\\TP_AS\D:\ASNT\D:	Finish
11	NT FAT	\\TP_AS\C:\PSP	Finish
12	NT FAT	\\TP_AS\C:	Finish
13	NT FAT	\\TP_AS\C:	Finish

Tape Information

Name: ELLENS FULL BACKUP
 ID: 172F Sequence No.: 1

Media Characteristics

Media Type:
 Density Code:
 Format Code: Volume Table

Format Information

First Format Date: 11/15/94 15:16:46
 Last Format Date: 11/15/94 15:16:46

Expiration Date: 11/15/95 23:59:58
 Location:
 Destroyed Date:

Host: TP_AS User: Administrator 11/20/94 4:28 PM

Viewing details about a tape

You can view additional details about a tape. This includes how many errors ARCserve has encountered while reading from or writing to the tape, how much data (total) has been written to the tape, and number of read and write passes ARCserve has made over the tape. In addition, you can see a list of the sessions that were created on the tape and a list of files backed up for each session.

Viewing statistical information about a tape

Use this record to determine if a tape should be retired because it has been in use for a long time or if you are getting a high number of errors.

To view statistical information about a tape:

1. Click the Statistics tab.

Category	Media Errors	Soft Read Errors	Soft Write Errors	MBytes Written	Usage Time	Times Formatted
Current Tape Statistics	0	0	21	209	0:28:17	
Total Tape Statistics	0	0	21	209	0:28:17	0

A media error means that some sort of data corruption has occurred on the tape and the read or write operation could not be successfully completed.

A soft error means that the tape drive detected a problem when it was trying to read from/write to the tape but was able to correct the problem by retrying the operation (or, in the case of a read, correcting the data through the use of its Error Correction Code). Use these fields to determine the "quality" of your tape. Remember, it is normal for a tape drive to always generate a certain number of soft read/write errors. The problem occurs when the number of errors is very high in relation to the amount of data being read/written.

Viewing information about what's on a tape

To view session and file information:

1. Highlight a tape.

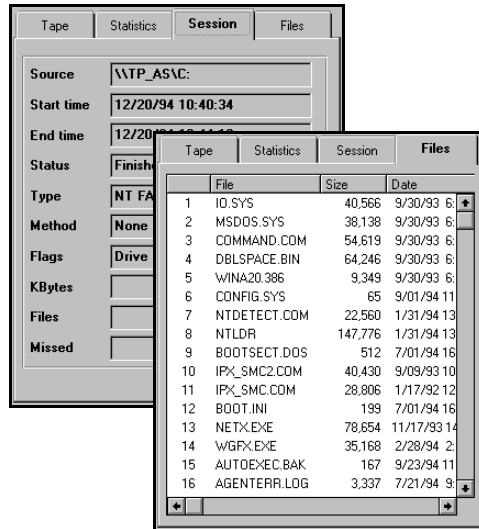
Session information is displayed in the *Sessions* section.

ID	Tape Name	Seq #	Last Format Date	Ses. #	Type	Path	Status
1	ELLENS FULL BACKUP [I/E	1	11/15/94 15:16	8	NT FAT	\\TWP_AS\VC:	Finish
2	ACCTG3 [D:F540]	1	11/03/94 9:37	9	NT NTFS	\\TWP_AS\D:	Finish
				10	ARCserve Dat.	\\TWP_AS\D:\ASNT\Dr	Finish
				11	NT FAT	\\TWP_AS\A:\PSP	Finish
				12	NT FAT	\\TWP_AS\VC:	Finish
				13	NT FAT	\\TWP_AS\VC:	Finish

2. Highlight the session you want to know more about.

3. Select the Session tab or the Files tab to see additional information about that session.


Select the Session tab to see information about the backup session or select the Files tab to see a listing of files that were backed up on this tape.




Keeping track of your tapes

The Tape View window includes a *Location* field that is initially empty but can be updated to help you keep track of where your backup tapes are located.

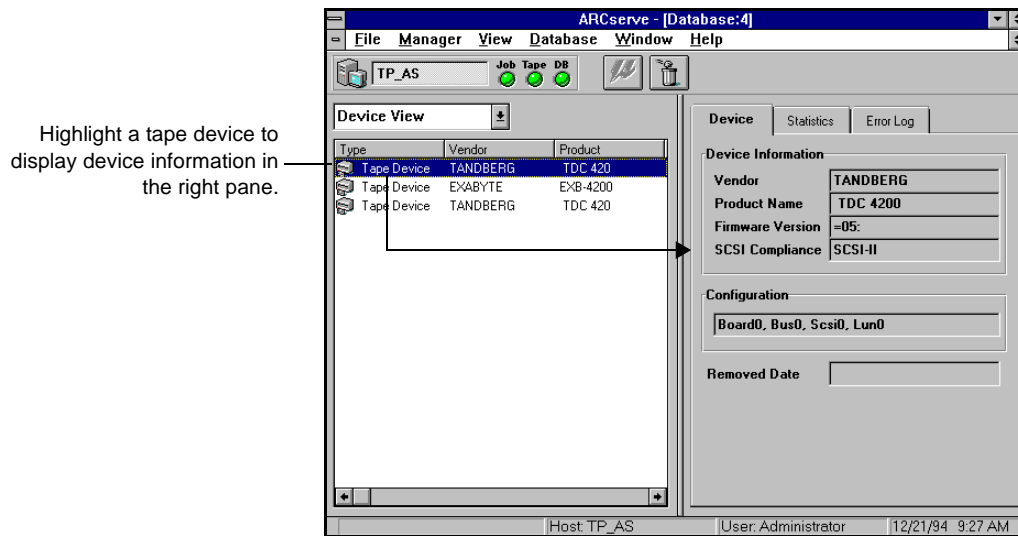
To enter a tape's location:

1. Highlight a tape.
2. Click the Modify button. 
3. Enter the tape's location.



Device View

The Device View window displays information about the tape devices used with ARCserve. This includes the type of device, vendor, SCSI ID, head clean count, and usage time.



Viewing
information about
a device

To view information about a tape device:

1. Highlight a tape device.

2. Select the Statistics tab or the Error Log tab.

A media error means that some sort of data corruption has occurred on the tape and the read or write operation could not be successfully completed.

A soft error means that the tape drive detected a problem when it was trying to write to the tape but was able to correct the problem by retrying the operation. Use this field to determine the "quality" of your tape. Remember, it is normal for a tape drive to always generate a certain number of soft write errors. The problem occurs when the number of errors is very high in relation to the amount of data being written.

Error Log tab

You will only see information in the Error Log if the tape drive has a critical error. An explanation of the fields in the Error Log appears below:

Field	Description
Code	Type of error - hardware or media (see explanation alongside the Statistics tab above).
Time	Time the error occurred.
Sense Info	SCSI error code.
Media	Number of media errors (see explanation above) that occurred during the job.
Soft Write	Number of soft write errors (see explanation above) that occurred during the job.
Soft Read	Number of soft read errors (see explanation above) that occurred during the job.
Tape Usage	Amount of time the tape was used during the job.
KBwritten	Amount of data written to the tape during the job.
Times Formatted	Number of times tape has been formatted.



11

Chapter

USING THE UTILITIES

ARCserve offers several utilities you can use to manage the files you have backed up and the files that are currently on your machine.

In this chapter, you will learn:

<i>Page</i>	
11-2	➤ What each utility is used for
11-3	➤ How to recover a database
11-6	➤ How to merge a tape
11-10	➤ How to scan a tape
11-13	➤ How to compare the contents of a tape
11-17	➤ How to count files on a machine
11-20	➤ How to purge files from a machine

Utilities basics

ARCserve offers several utilities that you can use to manage files. The utilities are:

- Recover Database - overwrites your ARCserve database files with a backed up copy of the database files.
- Merge Tape - adds information from tape into your ARCserve database.
- Scan Tape - reports information about one or more tape backup sessions.
- Compare - compares the contents of a tape to files on a machine.
- Count - counts files on a machine.
- Purge - deletes files from a machine.

This chapter describes these utilities and explains how and when to use them.

Using hidden menus in the Browser

ARCserve provides hidden menus in the Browser that give you additional options. These menus are accessed by clicking your right mouse button after highlighting a machine, tape group, or tape. The menu that you see depends on what you are highlighting.

If you are selecting a machine, the pop-up menu allows you to provide a user login name and password to access that machine.

If you are selecting a tape group or tape, the pop-up menu allows you to configure tape groups or perform maintenance on a tape. For more information about these options, refer to Chapter 9, “Managing Tapes & Tape Devices”.

Recover Database

This option allows you to restore database files from a tape to your ARCserve database.

Each time you run a backup job, ARCserve records information in its database about the machines, directories, and files that have been backed up, and the tapes 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 option lets you restore this information to your ARCserve database.

Unlike the Merge Tape option, which *appends* database information from tape to your existing database files, the Recover Database option *overwrites* existing database files.

Why would you need to use the Recover Database option?

You could use the Recover Database option anytime you lose your ARCserve database files, such as after a disaster like a fire, flood, or a machine crash, and you want to restore your files on a machine where you can create new ARCserve database files.

Heavy rains have caused the offices of your company to flood. All of the machines in your department are unusable. Luckily, your company has been able to find alternate space and rent computer equipment. You install ARCserve on your machine, and, with your department's backup tapes, you use the Recover Database option to recover the database. With the database recovered, you can now restore files to all of your department's rented machines.

Recovering a database

If you have some experience using ARCserve, you will be able to recover a database using these steps. If you need more information about using an ARCserve Manager window, refer to Chapter 3 “Basic Procedures”.



1. Click the Recover Database button from the Quick Access window.

Click to submit the recover job to the job queue. Click to indicate which database session you want to restore and to select options.

Use these fields to specify a tape that is not currently in your tape drive.

You can pick from a list of tapes currently in your tape drive.

Host: TP_AS User: Administrator 12/21/94 9:34 AM

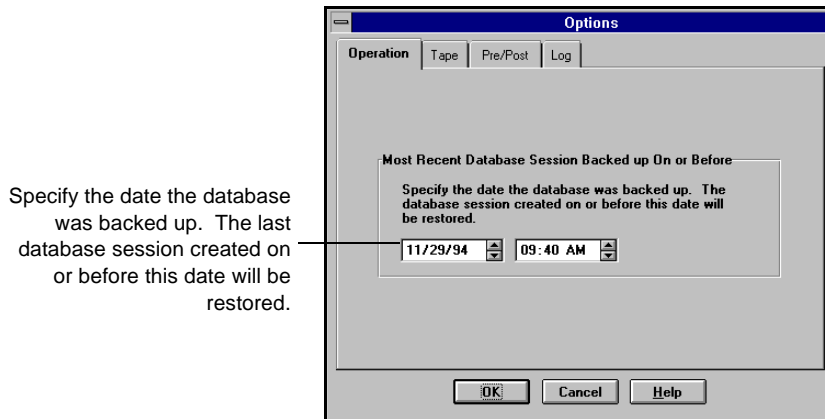
2. Specify which tape you want to use.

If the tape you want to use is not currently in a tape drive, you must type the information into the *Group* and *Tape* fields.

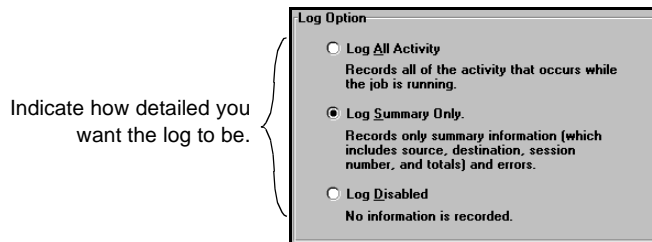
Enter the group and tape name. You must remember to mount this tape in your tape drive.


Group Tape

- Click the Options button  to indicate which database session you want to recover.



- Select the Log tab to determine how detailed you would like the Job Queue Log/Job View Database Log to be for this job.



- Run the job.
Click the Run/Schedule button  to run the job.

Merge Tape

This option allows you to take a tape that contains one or more backup sessions and add the information from the tape into your ARCserve database. The database information from the tape will be *appended* to your existing database files.

Each time you run a backup job, ARCserve records information in its database about the machines, directories, and files that have been backed up, and the tapes 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.

If you have a tape that has a backup session that is not included in the ARCserve database (the backup was created using ARCserve on a different machine, for example), you can use the Merge Tape option to get the tape's information into the database in ARCserve's home directory.

Unlike the Recover Database utility, which *overwrites* existing database files with database information from tape, the Merge Tape utility *appends* database information to your existing database files.

Why would you need to use the Merge Tape option?

You could use the Merge Tape utility if you need to restore files to an ARCserve machine that you did not use to create the backup.

The Chicago office of your company sends a tape to you at your New York office. The tape contains files that were backed up in Chicago using ARCserve. In order to restore the files you need, you use ARCserve's Merge Tape utility to add the database information from the tape to the database on your machine. You then restore the files you need to your machine.

You could also use the Merge Tape utility if you pruned (deleted) information from your ARCserve database that you now need. Refer to Chapter 14, “Administrating the ARCserve Server” for more information about pruning database records.

You need to restore a file that you created last January (eight months ago). You've saved all of the monthly backup tapes for the last 12 months, but your ARCserve system prunes database records that are older than six months.

In order to restore the file, you use ARCserve's Merge Tape utility to add the database information from the January tape to the database on your ARCserve machine.

Merging a tape into the ARCserve database

If you have some experience using ARCserve, you will be able to merge a tape using these steps. If you need more information about using an ARCserve Manager window, refer to Chapter 3, "Basic Procedures".



1. Click the Merge Tape button from the Quick Access window.

Click to submit the merge job to the queue. Click to specify options for this merge tape job. Specify which session to merge.

Use these fields to specify a tape that is not currently in your tape drive.

You can pick from a list of tapes currently in your tape drive.

Host: TP_AS User: Administrator 12/21/94 9:37 AM

2. Specify which tape you want to merge.

If the tape you want to use is not currently in a tape drive, you must type the information into the *Group* and *Tape* fields.

Enter the group and tape name. You must remember to mount this tape in your tape drive.

Group Tape


3. Indicate which session(s) you want to merge.



Enter a specific session to merge or click to merge all sessions on the tape.

If you do not know the session number, you can check the Activity Log or the Job View or Tape View windows in the database. If they are not available, you can use the Scan Tape utility to list the contents of the tape, including all session numbers.

4. Run the job.

Click the Run/Schedule button  to run the job.

Scan Tape

This option gives you information about your tape backup sessions. (Each source that you select to back up is saved on tape as an individual session.) You can scan a single session or an entire tape. Results of the Scan Tape job can be seen in the Job Queue or Job View Database Manager (by clicking the Log tab).

The Option button on the Scan Tape Manager allows you to select the level of detail you will see in the log. You can select to see just summary information and any errors that occurred during the scan or you can select to see each file that was backed up as well as the summary information and any errors. A sample report appears below:

In this example, session 1 on the tape ELLENS FULL BACKUP was scanned. The report lists each file that was backed up as well as summary information and any errors.

```
ARCserve for Windows NT -- Scan
Tape Name..... ELLENS FULL BACKUP, ID 172F
Sequence..... 1
Session..... 1
Workstation..... TP_AS
Path..... C:
Owner Name..... Administrator
Date & Time..... 12/09/94 12:44
Data Compression..... None
Session Type..... Windows NT FAT Drive
Session Method..... Backup
                  No Registry
                  Disk/Volume Backed Up
                  No Filters

C:
. <DIR> 0/00/80 0:00
IO.SYS 40,566 9/30/93 11:20
MSDOS.SYS 38,138 9/30/93 11:20
```

Why would you use the Scan Tape option?


Use the Scan Tape option to find out what is on a tape.

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 Tape option if you want a list of the files that were backed up.

Scanning a tape

If you already have some experience using ARCserve, you will be able to scan a tape using these steps. If you need more information about using an ARCserve Manager window, refer to Chapter 3, “Basic Procedures”.



1. Click the Scan Tape button from the Quick Access window.

The screenshot shows the ARCserve Scan Tape dialog box. Annotations include:

- Three arrows pointing to buttons at the top: "Click to submit the scan job to the job queue.", "Click to specify options for this scanning job.", and "Specify which session to scan."
- An arrow pointing to the "Source" fields: "Use these fields to specify a tape that is not currently in your tape drive."
- An arrow pointing to a list of tape groups: "You can pick from a list of tapes currently in your tape drive."

The dialog box contains the following fields and information:

- Host: TP_AS
- Job: Tape DB
- Source: Group [GROUP0] Tape [ELLENS FULL BACKUP] Sess. [1] All Sess.
- Device Information: Name [TANDBERG TDC 4200], Config. [Device:0 = Board:0, Bus:0, SCSI ID:0]
- Tape Information: Name [ELLENS FULL BACKUP], Sequence [1], Write Protect [No], ID [172F], Expiration Date [11/15/95 1]

11

2. Specify which tape you want to scan.

If the tape you want to use is not currently in a tape drive, you must type the information into the *Group* and *Tape* fields.

Enter the group and tape name. You must remember to mount this tape in your tape drive.

Group [GROUP0] Tape [REPORTS]

-
3. Indicate which session(s) you want to scan .



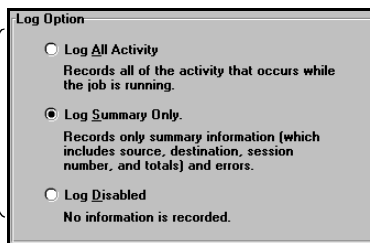
Sess. 1 All Sess.

Enter a specific session to scan or click to scan all sessions on the tape.

4. Determine how detailed you would like the Job Queue Log/Job View Database Log to be for this job .

Click the Options button  and select the Log tab.

Indicate how detailed you want the log to be.



Log Option

- Log All Activity
Records all of the activity that occurs while the job is running.
- Log Summary Only.
Records only summary information (which includes source, destination, session number, and totals) and errors.
- Log Disabled
No information is recorded.

5. Run the job.

Click the Run/Schedule button  to run the job.

Compare Tape

This option allows you to compare the contents of a tape session to files on a machine.

Results of the Compare Tape job can be seen in the Job Queue or Job View Database Manager (by clicking the Log tab).


The Option button on the Compare Tape Manager allows you to select the level of detail you will see in the log. You can select to see just summary information and any errors that occurred during the scan or you can create a report that lists each file that was compared and indicates if there is any difference in the file name, size, date, time, or attributes.

Why would you use the Compare Tape option?

You could use this option after a backup to verify that the backup copied all of the files to tape without error.

Comparing the contents of a tape

If you already have some experience using ARCserve, you will be able to compare a tape using these steps. If you need more information about using an ARCserve Manager window, refer to Chapter 3, “Basic Procedures”.



1. Click the Compare Tape button from the Quick Access window.

Click to submit the compare job to the job queue. Click to specify options for this compare job. Enter a session number to compare.

Click this tab to select the destination to compare against.

Use these fields to specify a tape that is not currently in your tape drive.

You can pick from a list of tapes currently in your tape drive.

Host: TP_AS User: Administrator 12/21/94 11:31 AM


2. Specify which tape you want to compare.

If the tape you want to use is not currently in a tape drive, you must type the information into the *Group* and *Tape* fields.

Enter the group and tape name. You must remember to mount this tape in your tape drive.

Group Tape

- Specify the session number you want to compare.
Enter the session number in the *Session* field.

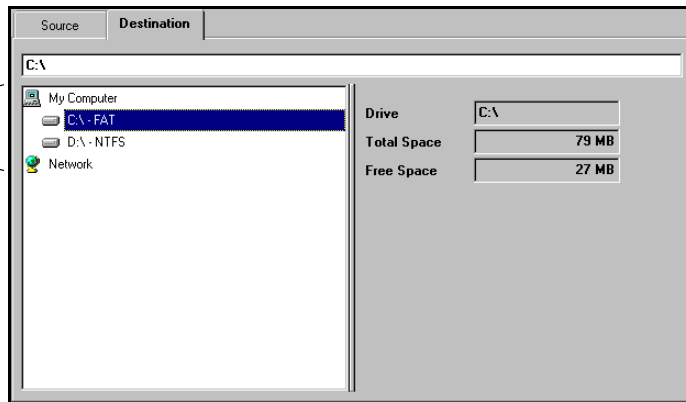


Enter a specific session number to compare. In this case, the first session on the tape will be compared.

If you do not know the session number, you can check the Activity Log or the Job View or Tape View windows in the database. If they are not available, you can use the Scan Tape utility to list the contents of the tape, including all session numbers.

- Select the Destination tab to specify what you want to compare against.

You can select a drive or directory on any machine. If you need help using the Browser, refer to Chapter 3, "Basic Procedures".



- Determine how detailed you would like the Job Queue Log/Job View Database Log to be for this job.


Click the Options button  and select the Log tab.

Indicate how detailed you want the log to be.

Log Option

- Log **A**ll Activity
Records all of the activity that occurs while the job is running.
- Log **S**ummary Only.
Records only summary information (which includes source, destination, session number, and totals) and errors.
- Log **D**isabled
No information is recorded.

6. Run the job.

Click the Run/Schedule button  to run the job.

Count

This option allows you to count the number of files and directories on a machine.

Results of the Count job can be seen in the Job Queue or Job View Database Manager (by clicking the Log tab).

The Options button on the Count Manager allows you to select the level of detail you will see in the log. You can select to see just totals and any errors that occurred or you can create a report that lists each file in every directory that was counted. This includes the file name, size, date, time, and file attributes. A sample report appears below:

In this example, all files on drive C: were counted. The report lists each file as well as its size, date, time, and file attributes.

```

ARCserve for Windows NT -- Count
Source..... C:
Start Time..... 10/17/94 18:34:20

IO.SYS                40,566    9/30/93 10:20  RHS-A
MSDOS.SYS             38,138    9/30/93 10:20  RHS-A
COMMAND.COM           54,619    9/30/93 10:20  ----A
DELSPACE.BIN          64,246    9/30/93 10:20  RHS-A
WINA20.386            9,349     9/30/93 10:20  ----A
CONFIG.SYS            65        9/01/94 15:44  ----A
NTDETECT.COM          22,560    1/31/94 17:28  ----A
NTLDR                 147,776   1/31/94 17:28  ----A
BOOTSECT.DOS          512       7/01/94 20:53  RHS-A
IPX_SMC2.COM          40,430    9/09/93 14:57  ----A
IPX_SMC.COM           28,806    1/17/92 16:29  ----A
BOOT.INI              199       7/01/94 20:55  R-S
NETX.EXE              78,654    11/17/93 18:14  ----A
WGFX.EXE              35,168    2/28/94  8:18  ----A
AUTOEXEC.BAT          167       9/23/94 15:57  ----A
AGENTERR.LOG          3,337     7/21/94 13:50  ----A

-----
C:\X
  <DIR>    10/05/94 16:14  ---D
  TRL.EXE 122,490    5/15/91  6:00  ----A
  TRL.HLP  43,371    5/15/91  6:00  ----A
  COPYRIGHT.SCR 1,388    5/15/91  6:00  ----A
  STI_SCI.EXE 21,569    5/15/91  6:00  ----A
  TRL.CFG   359     10/05/94 16:15  ----A
  VIEWFILE.TRL 512     10/05/94 16:15  ----A

Total Directories..... 29
Total Files..... 741
Total Skips..... 0
Total Size..... 39.38 MB
Elapsed Time..... 3s

Count Operation Successful.

```

Why would you use the Count option?

You could use this option to estimate how many tapes you will need for a backup.

Counting files

If you already have some experience using ARCserve, you will be able to count files using these steps. If you need more information about using an ARCserve Manager window, refer to Chapter 3, “Basic Procedures”.



1. Click the Count button from the Quick Access window .

Click to submit the count job to the job queue.

Click to determine how detailed you want the Activity Log for this job to be.

Click to filter the files and directories to be counted.


Select which files to count.

Drive	Total Space	Free Space
C:\	79 MB	27 MB

2. Specify what you want to count.

You can select one or more domains, machines, drives or directories.

3. Add any filters you may need.

You can add filters to include or exclude files or directories from being counted. To add filters, click the Filters  button.

For more information about filtering, refer to Chapter 7, “Customizing Your Jobs”.

4. Determine how detailed you would like the Job Queue Log/Job View Database Log to be for this job.


Click the Options button  and select the Log tab.

Indicate how detailed you want the log to be.

Log Option

- Log **A**ll Activity
Records all of the activity that occurs while the job is running.
- Log **S**ummary Only.
Records only summary information (which includes source, destination, session number, and totals) and errors.
- Log **D**isabled
No information is recorded.

5. Run the job.

Click the Run/Schedule button  to run the job.

Purge

This option allows you to delete files and directories from a machine. (Using filters, you can include or exclude files to be deleted.)

Results of the Purge job can be seen in the Job Queue or Job View Database Manager (by clicking the Log tab).

The Option button on the Purge Manager allows you to select the level of detail you will see in the log. You can select to see just totals and any errors that occurred or you can create a report that lists each file and directory that was deleted.

Why would you use the Purge option?

You could use this option to delete files that have not been used for a specific period of time.

Purging Files

If you already have some experience using ARCserve, you will be able to purge files using these steps. If you need more information about using an ARCserve Manager window, refer to Chapter 3, “Basic Procedures”.



1. Click the Purge button from the Quick Access window.

Click to submit the purge job to the job queue.

Click to determine how detailed you want the Activity Log for this job to be.


Click to filter the files and directories to be purged.

Select which files to purge.

2. Specify what you want to purge.


You can select a drive, directory, or specific files on a machine.

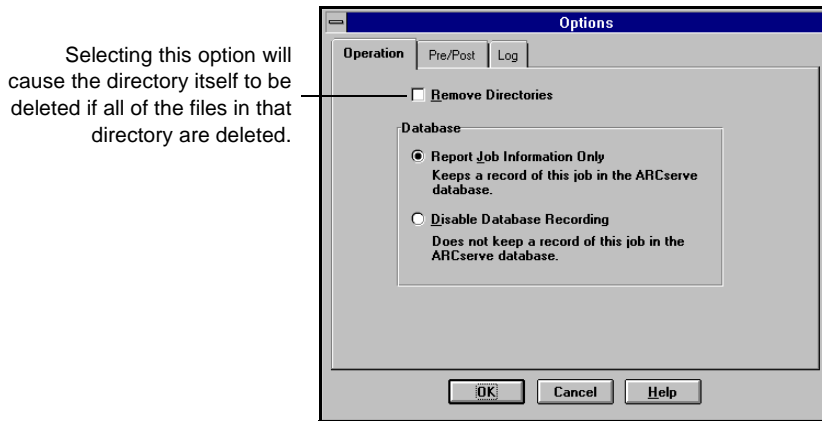
3. Add any filters you may need.

You can add filters to include or exclude files or directories from being purged. To add filters, click the Filter  button.

For more information about filtering, refer to Chapter 7 “Customizing Your Jobs”.

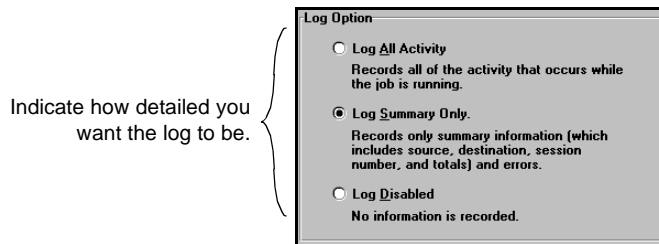
-
4. Indicate if you would like to delete directories that have been completely purged.

Click the Options button  and select *Remove Directories*.




5. Determine how detailed you would like the Job Queue Log/Job View Database Log to be for this job.

Click the Options button  and select the Log tab.



6. Run the job.

Click the Run/Schedule button  to run the job.



Be certain that any drive, directory, or files you have selected are really the ones you want to purge!

12

Chapter

VIEWING AND PRINTING REPORTS

This chapter discusses the various logs and reports that are available in ARCserve.

In this chapter, you will learn:

<i>Page</i>	
12-3	➤ About the Activity Log
12-6	➤ What reports are available in the Reports Manager
12-14	➤ How to determine what you will see in the Job Queue or Job View Database Log

Report basics

ARCserve has several features that can help you sort through all of the information available about your ARCserve system and the jobs you have run.

Quick Status Manager

The Quick Status Manager allows you to view ARCserve's Activity Log. This log provides a journal of activity within ARCserve.

Reports Manager

The Reports Manager offers several reports that contain information from ARCserve's database. These reports provide you with a convenient way of getting a hard copy printout of database information. The following reports are available:

- Job Report
- Tape Media Report
- Tape Session Report
- Tape Session Detail Report
- Tape Device Report

Job Queue Manager and Job View Database Manager

The Job Queue Manager and Job View Database Manager offer log information for all executed jobs. The level of detail in this report is determined when you set up a job.

The remainder of this chapter explains how to use the Quick Status Manager and the Reports Manager, provides a description of each log and report, and explains how to determine the level of detail you will see in the Job Queue Summary or Job View Database Log.

Viewing 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
- Any errors and warnings that occurred

Why would you need to view this log?

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.

Accessing the Activity Log






To access the Activity Log:

1. Click the Quick Status button from the Quick Access window.

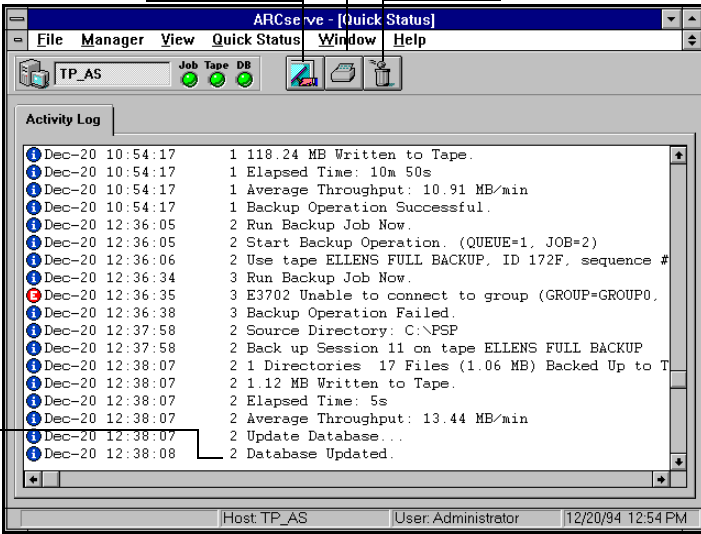
The Activity Log is displayed:

Click to update the information in the log. Click to print the log. Click to clear the log.

Each line has an icon that rates the severity of the message:

-  = Informational
-  = Warning
-  = Error

This is the Job ID. (Each job has a unique Job ID.)




Depending upon how your system is configured, you may also see a Tape Engine Log in the Quick Status Manager. This log displays debugging information about the Tape Engine. Refer to Chapter 14, "Administrating the ARCserve Server" for more information about displaying the Tape Engine Log.

Using the hidden menu

If you click your right mouse button in the Activity Log, the following menu appears:

Use to change the font that the log is displayed in. — **F**onts...

Use to print the log. — **P**rint...

Use to search for specific text within the log. — **F**ind... Use to repeat the text search.

Find Next F3

Refresh F5 Use to update information in the log..

Printing the Activity Log

You can print the Activity Log.

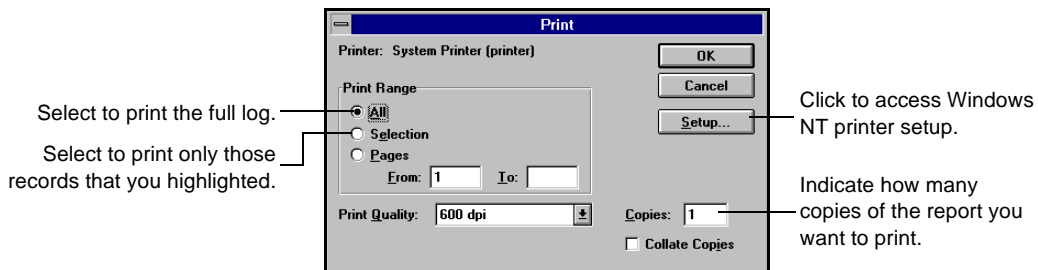
Follow these steps to print the log:

1. Select the records you want to print.
If you only want to print selected records, highlight the first record that you want, press and hold down your Shift key while you select the remaining records.

2. Click the Print button. 


You can also click your right mouse button.

3. Select the print options you want.




4. Click OK.

Refreshing the Activity Log

Click the Refresh button  if you have been looking at the Activity Log for awhile and you want to update the information in the log without having to close the Manager window and then re-open it. You can then press the END key to move to the end of the log to see the latest information.

Deleting the Activity Log

Click the Delete button  to delete all of the information from the Activity Log.

Using the Reports Manager

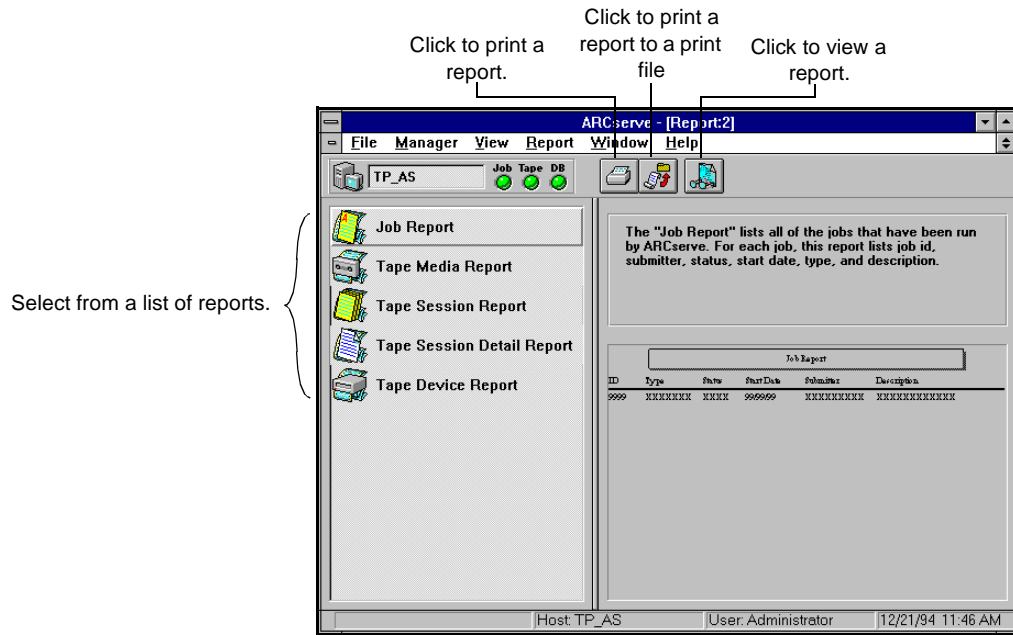
The Reports Manager offers several reports that contain information from ARCserve's database. To open the window:



1. Click the Reports button from the Quick Access window .

The ARCserve Reports Manager window appears.

The Reports Manager window displays a list of the reports that are available.



Printing a report You can print any report that is listed in the Reports Manager.

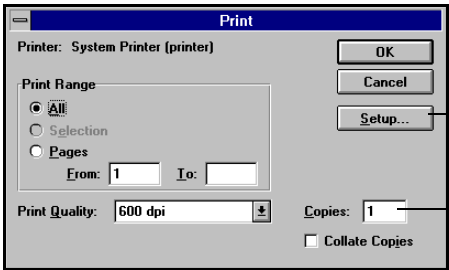
Follow these steps to print a report:

1. Highlight the report you want to print.
You can also print a report that you are currently viewing.

2. Click the Print to Printer button.



3. Select the print options you want.



Specify which pages to print.

Click to access Windows NT printer setup.

Indicate how many copies of the report you want to print.

The 'Print' dialog box shows 'Printer: System Printer (printer)'. Under 'Print Range', 'All' is selected. 'From: 1' and 'To:' are visible. 'Print Quality' is set to '600 dpi'. 'Copies: 1' is indicated, and 'Collate Copies' is unchecked. Buttons for 'OK', 'Cancel', and 'Setup...' are present.

4. Click OK.

Printing a report to a file

You can print any report listed in the Reports Manager to an ASCII file.

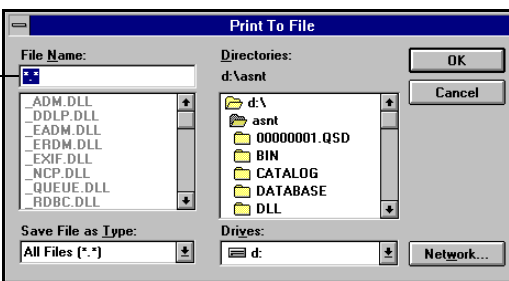
Follow these steps to print a report to a print file:

1. Highlight the report you want to print.

2. Click the Print to File button.



3. Enter a name for the report file.




Enter the name for the report file. You may also need to specify the directory or drive.


The 'Print To File' dialog box has a 'File Name:' field with a blue asterisk icon. Below it is a list of files including '_ADM.DLL', '_DDL.DLL', '_EADL.DLL', '_ERDM.DLL', '_EXIF.DLL', '_NCP.DLL', '_QUEUE.DLL', and '_RDBC.DLL'. The 'Directories:' section shows 'd:\asnt' with subdirectories 'asnt', '00000001.QSD', 'BIN', 'CATALOG', 'DATABASE', and 'DLL'. The 'Save File as Type:' is set to 'All Files (*.*)'. The 'Drives:' section shows 'd:' selected. Buttons for 'OK', 'Cancel', and 'Network...' are present.

4. Click OK.

Viewing a report Follow these steps to view a report:

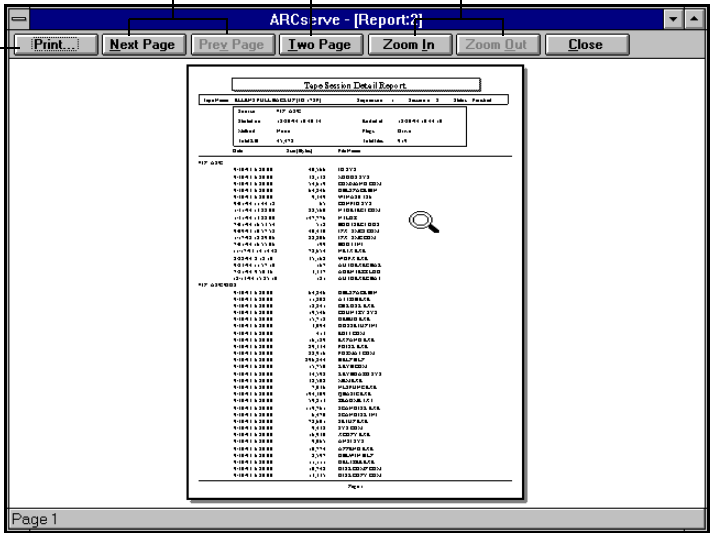
1. Highlight the report you want to view .
2. Click the Print Preview button. 

You can also double-click on the report name.

3. If applicable, select the tape, session, or device you want to view information about.
4. Click your cursor, which now looks like , to enlarge the text on the report.

Click to move to the next or previous page of this report. Click to view two pages of this report at once. Click to enlarge/reduce the text on the report.

Click to print the report.



Device	TAP 030E	Session #	133861.00 04 10
Label	Pages	Page	Block
LENGTH	13,573	140000	1.15
File	File Name	File Type	
030E1.00 00	030E1.00 00	030E1.00 00	030E1.00 00
030E1.00 01	030E1.00 01	030E1.00 01	030E1.00 01
030E1.00 02	030E1.00 02	030E1.00 02	030E1.00 02
030E1.00 03	030E1.00 03	030E1.00 03	030E1.00 03
030E1.00 04	030E1.00 04	030E1.00 04	030E1.00 04
030E1.00 05	030E1.00 05	030E1.00 05	030E1.00 05
030E1.00 06	030E1.00 06	030E1.00 06	030E1.00 06
030E1.00 07	030E1.00 07	030E1.00 07	030E1.00 07
030E1.00 08	030E1.00 08	030E1.00 08	030E1.00 08
030E1.00 09	030E1.00 09	030E1.00 09	030E1.00 09
030E1.00 10	030E1.00 10	030E1.00 10	030E1.00 10
030E1.00 11	030E1.00 11	030E1.00 11	030E1.00 11
030E1.00 12	030E1.00 12	030E1.00 12	030E1.00 12
030E1.00 13	030E1.00 13	030E1.00 13	030E1.00 13
030E1.00 14	030E1.00 14	030E1.00 14	030E1.00 14
030E1.00 15	030E1.00 15	030E1.00 15	030E1.00 15
030E1.00 16	030E1.00 16	030E1.00 16	030E1.00 16
030E1.00 17	030E1.00 17	030E1.00 17	030E1.00 17
030E1.00 18	030E1.00 18	030E1.00 18	030E1.00 18
030E1.00 19	030E1.00 19	030E1.00 19	030E1.00 19
030E1.00 20	030E1.00 20	030E1.00 20	030E1.00 20
030E1.00 21	030E1.00 21	030E1.00 21	030E1.00 21
030E1.00 22	030E1.00 22	030E1.00 22	030E1.00 22
030E1.00 23	030E1.00 23	030E1.00 23	030E1.00 23
030E1.00 24	030E1.00 24	030E1.00 24	030E1.00 24
030E1.00 25	030E1.00 25	030E1.00 25	030E1.00 25
030E1.00 26	030E1.00 26	030E1.00 26	030E1.00 26
030E1.00 27	030E1.00 27	030E1.00 27	030E1.00 27
030E1.00 28	030E1.00 28	030E1.00 28	030E1.00 28
030E1.00 29	030E1.00 29	030E1.00 29	030E1.00 29
030E1.00 30	030E1.00 30	030E1.00 30	030E1.00 30
030E1.00 31	030E1.00 31	030E1.00 31	030E1.00 31
030E1.00 32	030E1.00 32	030E1.00 32	030E1.00 32
030E1.00 33	030E1.00 33	030E1.00 33	030E1.00 33
030E1.00 34	030E1.00 34	030E1.00 34	030E1.00 34
030E1.00 35	030E1.00 35	030E1.00 35	030E1.00 35
030E1.00 36	030E1.00 36	030E1.00 36	030E1.00 36
030E1.00 37	030E1.00 37	030E1.00 37	030E1.00 37
030E1.00 38	030E1.00 38	030E1.00 38	030E1.00 38
030E1.00 39	030E1.00 39	030E1.00 39	030E1.00 39
030E1.00 40	030E1.00 40	030E1.00 40	030E1.00 40
030E1.00 41	030E1.00 41	030E1.00 41	030E1.00 41
030E1.00 42	030E1.00 42	030E1.00 42	030E1.00 42
030E1.00 43	030E1.00 43	030E1.00 43	030E1.00 43
030E1.00 44	030E1.00 44	030E1.00 44	030E1.00 44
030E1.00 45	030E1.00 45	030E1.00 45	030E1.00 45
030E1.00 46	030E1.00 46	030E1.00 46	030E1.00 46
030E1.00 47	030E1.00 47	030E1.00 47	030E1.00 47
030E1.00 48	030E1.00 48	030E1.00 48	030E1.00 48
030E1.00 49	030E1.00 49	030E1.00 49	030E1.00 49
030E1.00 50	030E1.00 50	030E1.00 50	030E1.00 50
030E1.00 51	030E1.00 51	030E1.00 51	030E1.00 51
030E1.00 52	030E1.00 52	030E1.00 52	030E1.00 52
030E1.00 53	030E1.00 53	030E1.00 53	030E1.00 53
030E1.00 54	030E1.00 54	030E1.00 54	030E1.00 54
030E1.00 55	030E1.00 55	030E1.00 55	030E1.00 55
030E1.00 56	030E1.00 56	030E1.00 56	030E1.00 56
030E1.00 57	030E1.00 57	030E1.00 57	030E1.00 57
030E1.00 58	030E1.00 58	030E1.00 58	030E1.00 58
030E1.00 59	030E1.00 59	030E1.00 59	030E1.00 59
030E1.00 60	030E1.00 60	030E1.00 60	030E1.00 60
030E1.00 61	030E1.00 61	030E1.00 61	030E1.00 61
030E1.00 62	030E1.00 62	030E1.00 62	030E1.00 62
030E1.00 63	030E1.00 63	030E1.00 63	030E1.00 63
030E1.00 64	030E1.00 64	030E1.00 64	030E1.00 64
030E1.00 65	030E1.00 65	030E1.00 65	030E1.00 65
030E1.00 66	030E1.00 66	030E1.00 66	030E1.00 66
030E1.00 67	030E1.00 67	030E1.00 67	030E1.00 67
030E1.00 68	030E1.00 68	030E1.00 68	030E1.00 68
030E1.00 69	030E1.00 69	030E1.00 69	030E1.00 69
030E1.00 70	030E1.00 70	030E1.00 70	030E1.00 70
030E1.00 71	030E1.00 71	030E1.00 71	030E1.00 71
030E1.00 72	030E1.00 72	030E1.00 72	030E1.00 72
030E1.00 73	030E1.00 73	030E1.00 73	030E1.00 73
030E1.00 74	030E1.00 74	030E1.00 74	030E1.00 74
030E1.00 75	030E1.00 75	030E1.00 75	030E1.00 75
030E1.00 76	030E1.00 76	030E1.00 76	030E1.00 76
030E1.00 77	030E1.00 77	030E1.00 77	030E1.00 77
030E1.00 78	030E1.00 78	030E1.00 78	030E1.00 78
030E1.00 79	030E1.00 79	030E1.00 79	030E1.00 79
030E1.00 80	030E1.00 80	030E1.00 80	030E1.00 80
030E1.00 81	030E1.00 81	030E1.00 81	030E1.00 81
030E1.00 82	030E1.00 82	030E1.00 82	030E1.00 82
030E1.00 83	030E1.00 83	030E1.00 83	030E1.00 83
030E1.00 84	030E1.00 84	030E1.00 84	030E1.00 84
030E1.00 85	030E1.00 85	030E1.00 85	030E1.00 85
030E1.00 86	030E1.00 86	030E1.00 86	030E1.00 86
030E1.00 87	030E1.00 87	030E1.00 87	030E1.00 87
030E1.00 88	030E1.00 88	030E1.00 88	030E1.00 88
030E1.00 89	030E1.00 89	030E1.00 89	030E1.00 89
030E1.00 90	030E1.00 90	030E1.00 90	030E1.00 90
030E1.00 91	030E1.00 91	030E1.00 91	030E1.00 91
030E1.00 92	030E1.00 92	030E1.00 92	030E1.00 92
030E1.00 93	030E1.00 93	030E1.00 93	030E1.00 93
030E1.00 94	030E1.00 94	030E1.00 94	030E1.00 94
030E1.00 95	030E1.00 95	030E1.00 95	030E1.00 95
030E1.00 96	030E1.00 96	030E1.00 96	030E1.00 96
030E1.00 97	030E1.00 97	030E1.00 97	030E1.00 97
030E1.00 98	030E1.00 98	030E1.00 98	030E1.00 98
030E1.00 99	030E1.00 99	030E1.00 99	030E1.00 99
030E1.00 100	030E1.00 100	030E1.00 100	030E1.00 100

Job Report

This report provides a brief listing of all jobs that have been run by ARCserve. It contains the information that is found in the Database Manager on the Job View window. Each time ARCserve runs a job, a record is added to ARCserve's database. This record contains basic information about the job. When the job is completed, the record is updated to reflect the status and completion time of the job.

For each job, this report includes the following:

- Job ID
- Job type
- Job status
- Start time
- Submitter
- Description

Why would you need to view this report?

You can view the report to see what jobs have been run the past few days.

Tape Media Report

This report provides information about the tapes you use with ARCserve. It contains the information that is found in the Database Manager on the Tape View window. This report gives you a convenient way of getting a hard copy printout of your Tape View information.

When you go to view this report you will be asked to select a tape.

This report includes the following:

- Tape name, ID, and sequence number
- Media type
- Usage time
- Creation, format, and expiration date
- Amount of data written to the tape
- Number of errors
- Tape usage statistics

Why would you need to view this report?

You may need to view this report if you are experiencing problems with a tape. This report can help you decide if it is time to “retire” the tape.

Tape Session Report

This report provides information about all of the backup sessions that are on tape. (Each source you select to back up is saved on tape as an individual session.) The report contains the session information that is found in the Database Manager on the Job View or Tape View window when you highlight a record and click the *Session* tab.

When you go to view this report you will be asked if you want the report to include all tapes or just one that you select.

For each tape session, this report includes the following:

- Session number
- Status
- Source selected
- Start time
- End time
- Session method
- Flags
- Total files/KB in the session

You can view this report for all tapes or just one.

Why would you need to view this report?

You can view this report to get a general idea of what is backed up on tape when you need to restore files.

Tape Session Detail Report

This report includes all of the information found in the Tape Session Report and also lists every file that was backed up in each session. The report contains the session information that is found in the Database Manager on the Job View or Tape View windows when you highlight a record and click the *Files* tab.

When you go to view this report you will be asked to select a tape and a session from that tape.

Why would you need to view this report?

You can view this report if you need to know which files are backed up in each backup session. This is useful when you need to restore files.

Tape Device Report

This report provides information about all of the tape devices you are using with ARCserve. It contains the information that is found in the Database Manager on the Device View window. This report gives you a convenient way of getting a hard copy printout of your Device View information.

When you go to view this report you will be asked to select a tape device.

For each tape device, this report includes the following:

- Vendor
- Firmware revision
- SCSI ID
- Device type
- Last head cleaning date
- Head cleaning count
- Amount of usage
- Errors

Why would you need to view this report?

You may need to view this report if you are experiencing problems with a tape device. It can indicate if it is time to clean the heads, or if you have the wrong firmware installed.

Job Queue Summary Log/Job View Database Log

The Job Queue Manager and Job View Database Manager offer log information for all executed jobs. The level of detail in this report is determined when you set up a job. You can select 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). You also have the ability to disable reporting.

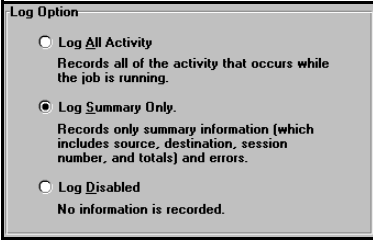
To select the level of detail:

1. Click the Options button. 

This button is available on the Backup, Copy, Restore, and the utilities Managers.

2. Click the Log tab.

Indicate how detailed you want the log to be.



The image shows a dialog box titled "Log Option" with three radio button options. The first option is "Log All Activity" with the description "Records all of the activity that occurs while the job is running." The second option is "Log Summary Only" with the description "Records only summary information (which includes source, destination, session number, and totals) and errors." The third option is "Log Disabled" with the description "No information is recorded." The "Log Summary Only" option is selected.

Option	Description
<input type="radio"/> Log A ll Activity	Records all of the activity that occurs while the job is running.
<input checked="" type="radio"/> Log S ummary Only.	Records only summary information (which includes source, destination, session number, and totals) and errors.
<input type="radio"/> Log D isabled	No information is recorded.

3. Select the option you want.

13

Chapter

DISASTER RECOVERY

This chapter provides you with information about recovering data in the event of a system failure.

In this chapter, you will learn:

Page

- | | | |
|------|---|---|
| 13-2 | ➤ | Tips to help make recovering from disaster easier |
| 13-4 | ➤ | How to recover crashed Windows NT machines |

Disaster recovery basics

This chapter describes how you should use ARCserve to best protect and recover your data in the event of a system failure. Refer to Chapter 5, “Restoring Your Data” for details about the restore process.

Make sure you
back up
ARCserve’s
database often

Whenever you select your ARCserve home directory for backup, ARCserve’s database is backed up as well. Always maintain a current backup of your ARCserve database files so that recovering from a disaster is easier.

Always maintain
current backups

The most important thing you can do to ensure against data loss is to maintain current backups of all your Windows NT machines. 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 tape rotation scheme using different backup methods and scheduling as a means of maintaining current backups.

If something happens to any machine between full backups, you can recover with the last full backup (a backup job with the *Full (Clear Archive Bit)* method selected), and your most recent differential or incremental backup (which contains all the modifications to the data since the last full backup). Refer to Chapter 4, “Backing Up to Tape”, for details on backup methods.

The following books will be useful in the following sections should you need to use ARCserve to recover a Windows NT machine:

- Windows NT Manuals - If you need to install the Windows NT operating system.

In addition, you should have the appropriate installation disks and the backup tapes to restore from.

Follow the instructions included in the next few sections to recover your Windows NT machines.

Recovering a crashed Windows NT machine

To recover a Windows NT machine, follow the instructions outlined below:

1. Re-install the Windows NT operating system .

Refer to your Windows NT documentation to re-install the operating system.

Refer to Chapter 1, "Getting Started", to re-install ARCserve

2. (Optional) Re-install ARCserve.

If the machine that crashed was an ARCserve machine, you will need to re-install ARCserve before you can proceed to recover *any* machine on your Microsoft network.

An ARCserve machine is a Windows NT machine that has ARCserve for Windows NT installed on it. In addition to the ARCserve software, it contains a copy of the ARCserve database, which will need to be recovered in order to recover the machine.

Refer to Chapter 11, "Using the Utilities", for details on how to use the Recover Database utility

3. (Optional) Restore the most current ARCserve database to the ARCserve machine.

Again, if the machine that crashed was your ARCserve machine, you will need to recover the latest version of your database back to the ARCserve machine. Use the Recover Database utility to perform this step.

The ARCserve database contains information about what has been previously backed up. You will need to recover the ARCserve database so that you can restore the rest of the information on the host.

Refer to Chapter 5, "Restoring Your Data" for details on restoring

4. Use the "Restore by Tree" source view in the Restore Manager to select what data you want to recover .

Now that you have restored the database, use the “Restore by Tree” source view to display the data that was backed up from each machine. You can pick and choose what you want to recover (or select the entire machine).

5. Select the machine you want to recover.

Click the Destination tab in the Restore Manager and select the machine you want to recover the data to.

6. Click the Options button to specify destination options.



If you are trying to restore the machine to the state it was in before it crashed, then you should:

- Select a Destination option that overwrites any conflicting files. If you are unsure, you can choose to rename any files that exist on the machine.
- Select the *Restore Registry* option.

Refer to Chapter 5, “Restoring Your Data”, for details on these restore options.

7. Click the Run/Schedule button to run the job.

ARCserve will then tell you the tapes that you will need to recover the machine.



14

Chapter

ADMINISTRATING THE ARCSERVE SERVER

The ARCServe Server consists of the Job Engine, Tape Engine, and Database Engine. The ARCServe Server Admin is your gateway to maintaining and managing them.

In this chapter, you will learn:

Page

- | | | |
|-------|---|---|
| 14-6 | ➤ | How to view information about an individual engine |
| 14-9 | ➤ | How to start, stop, pause, and resume functionality of the ARCServe engines from the ARCServe Server Admin dialog box |
| 14-11 | ➤ | How to configure the ARCServe engines |

The ARCserve Server and the ARCserve Server Admin

What is the ARCserve Server?

The ARCserve Server consists of three functional engines. They are the Job Engine, Tape Engine, and Database Engine. For a description of these engines and how they work together to make up the ARCserve Server, refer to Chapter 2, "About ARCserve for Windows NT".

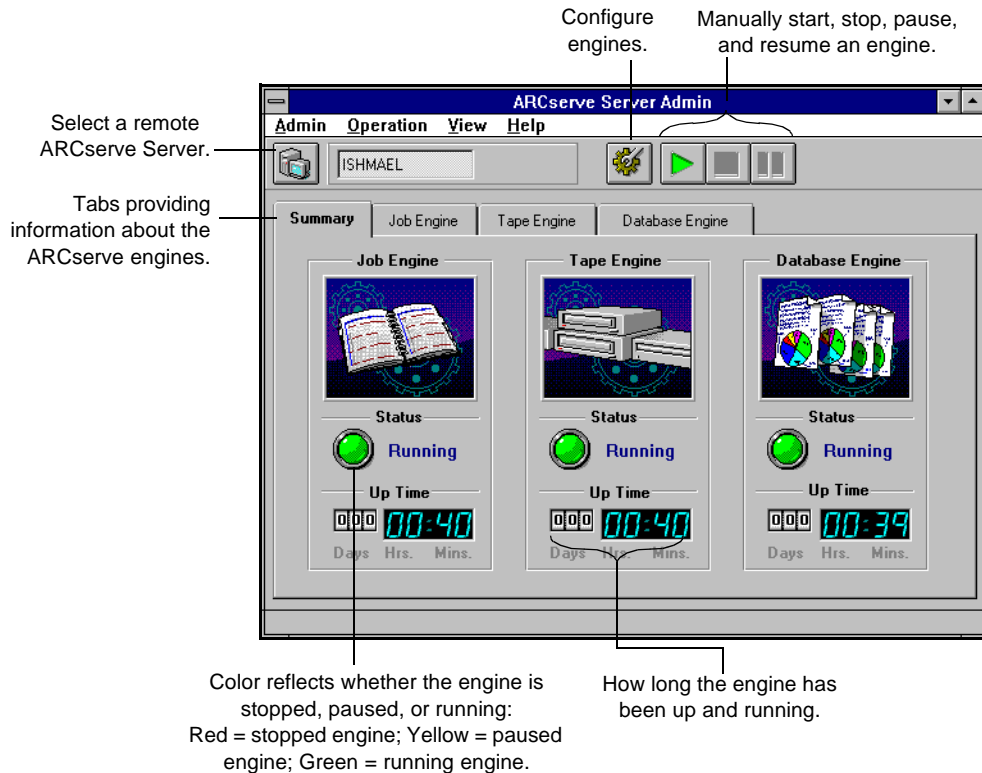
Getting to the ARCserve Server Admin

The ARCserve Server Admin is your gateway to the ARCserve Server. To access it:



1. In your ARCserve for Windows NT program group, double-click the ARCserve Server Admin icon.

The ARCserve Server Admin appears, as shown below:



Remote Administration

You can manage an ARCserve Server from a remote machine. Refer to Chapter 1, “Getting Started” to learn how install the ARCserve Server Administrator on a remote machine.

If you are running the ARCserve Server Administrator remotely, you can select a server by clicking the Select Server button. Refer to Chapter 3, “Basic Procedures” for information on the Select Server button.

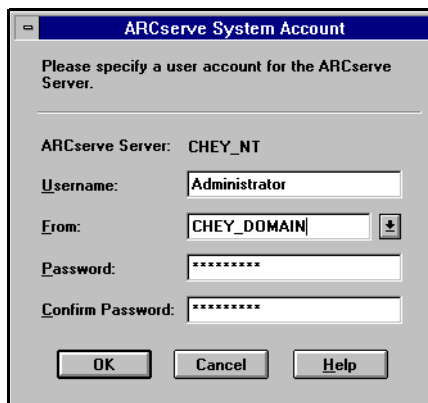
Changing the ARCserve System Account

The ARCserve Server requires a valid user account on the host Windows NT machine. You must enter a system account during the installation program. You can change the system account at any time in the ARCserver Server Administrator.

To change the ARCserve system account:

1. Choose ARCserve System Account from the File menu .

The ARCserve System Account dialog box opens:



ARCserve System Account

Please specify a user account for the ARCserve Server.

ARCserve Server: CHEY_NT

Username: Administrator

From: CHEY_DOMAIN

Password: *****

Confirm Password: *****

OK Cancel Help

Enter a valid username and password for the Windows NT machine on which the ARCserve Server is installed.

2. Click OK.

The ARCserve System Account is updated.


Getting information about the ARCserve engines

The ARCserve Server Admin has four tabs, each showing you various information.

- Summary tab - shows you the current state of all the ARCserve engines.
- Job Engine tab - displays information about jobs, such as the number of active and pending jobs.
- Tape Engine tab - displays information about jobs using the tape engine, such as the type of job it is and who submitted it.
- Database Engine tab - displays information about database pruning.

Summary information about all the ARCserve engines

At first glance, the ARCserve Server Admin shows you the current state of the various ARCserve engines (whether each is running, paused, or stopped). It also shows you how long each has been running.



Viewing information about an individual engine

Job Engine information

To get specific information about the ARCserve Job Engine:

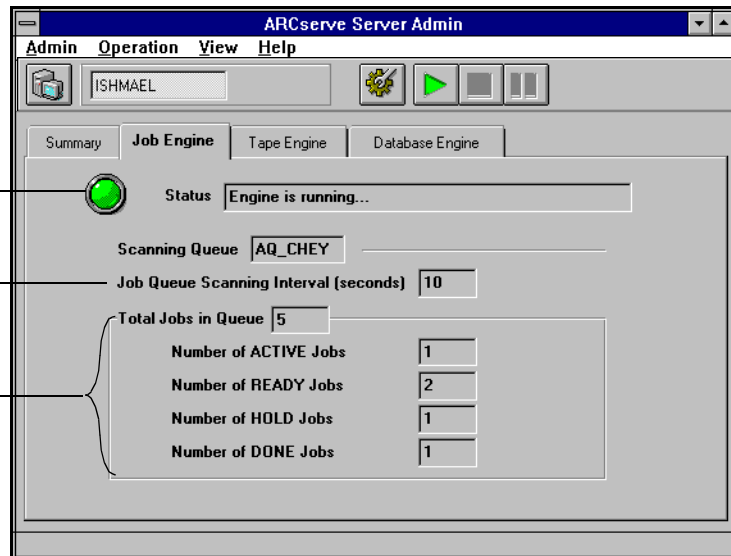
1. From the ARCserve Server Admin, select the Job Engine tab.

The ARCserve Server Admin displays information about the Job Engine, as seen below:

Describes the state of the engine, such as if it is running, paused, stopped, etc.

Tells you how often the job queue is scanned for jobs ready to be executed. This information is set in the Job Engine configuration.

This information is taken from the job queue. It shows the number of jobs currently in the queue, and then separates them by job status.



Tape Engine information

To view information specific to the ARCserve Tape Engine:

1. From the ARCserve Server Admin, select the Tape Engine tab.

Details specific to the Tape Engine appear, similar to the screen shown below:

The screenshot shows the 'ARCserve Server Admin' window with the 'Tape Engine' tab selected. The status is 'Engine is running...' indicated by a green light. Below this, there are two tables: 'Total Active Tape Engine Jobs: 2' and 'Total Active Data Transfer Groups: 0'. The first table lists jobs for 'ARCserve Manager' and 'audrey'. The second table is empty.

Indicates whether the Tape Engine is running or stopped.

The number of active jobs currently using the Tape Engine...

...and some information about these active tape groups, such as the group name, tape name, etc.

Client Name	Job Type	Process State
ARCserve Manager	Device Management	None
audrey	Read/Write	None

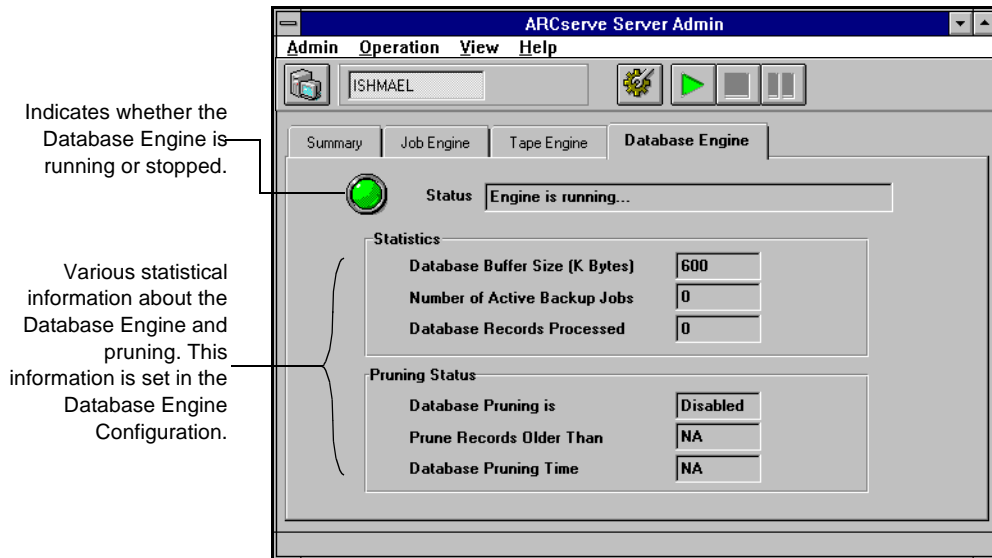
Group Name	Device	Tape Name	Tape ID	Seq	WProt	KB Used
------------	--------	-----------	---------	-----	-------	---------

Database Engine information

To view information about the Database Engine:

1. From the ARCserve Server Admin, select the Database Engine tab.

Information about the Database Engine is displayed:



Viewing and changing the running status of the engines

Although you are able to view and change the running status of the ARCserve Server engines via the ARCserve Manager windows, the ARCserve Server Admin can be used to do the same.

What happens when an engine is stopped


A stopped engine is an engine that is completely off-line. This may be because you have not yet started it, you brought it down manually for some other purpose, or, in the course of running, it has encountered an error and shut itself down. Whatever the reason, it means that the services that the engine provides are not available.

ARCserve's engines are designed to run independently of each other. That means that if, say, you stop the Tape Engine, the Database Engine and the Job Engine are not affected. They continue to run, providing and 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 requiring a tape drive is encountered, the job is launched by the Job Engine as usual, but the job will fail because the Tape Engine is not there to communicate with the tape drives. The Database Engine then logs this information.



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

Starting an engine from the ARCserve Server Admin


1. In the ARCserve Server Admin, select the tab for the engine you want to start.
2. Click the Start Engine button. 

You can verify that the engine is started by checking the *Status* field (seen in the engine's tab) or by checking for a green status light under the engine's name (seen in the Summary tab).

Stopping an engine from the ARCserve Server Admin

Follow the steps below to completely halt services from an engine:

1. In the ARCserve Server Admin, select the tab for the engine you want to stop.

2. Click the Stop Engine button. 

You can verify that the engine is stopped by checking the engine's *Status* field (in the engine's tab) or by checking for a red status light under the engine's name (in the Summary tab).


Pausing the Job Engine from the ARCserve Server Admin



Follow the steps below to temporarily pause the Job Engine.

Only the Job Engine can be paused.

1. In the ARCserve Server Admin, select the Job Engine tab.

2. Click the Pause Engine button. 

You can verify that the Job Engine is paused by checking the *Status* field in the Job Engine tab, or by checking for a yellow status light under Job Engine in the Summary tab.

Configuring the ARCserve engines

The ARCserve Server Admin allows you to configure each engine to suit your particular needs. The following sections show you how to configure the ARCserve engines.

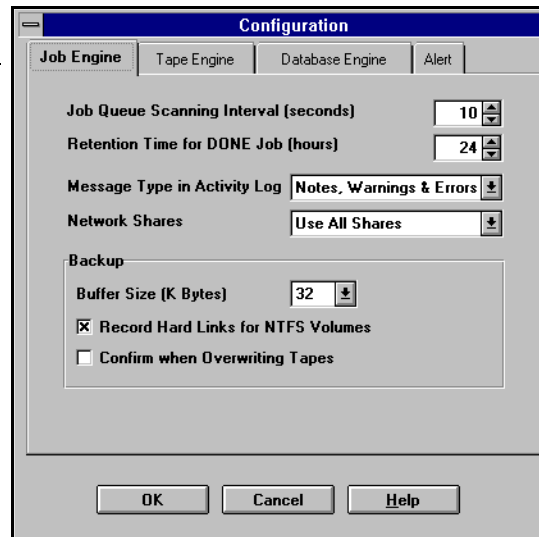
1. From the ARCserve Server Admin, click the Configure

Engine button. 

The Configuration dialog box appears. By default, the Job Engine tab is selected:

Select the tab for the engine you want to configure. By default, the Job Engine tab is selected first.

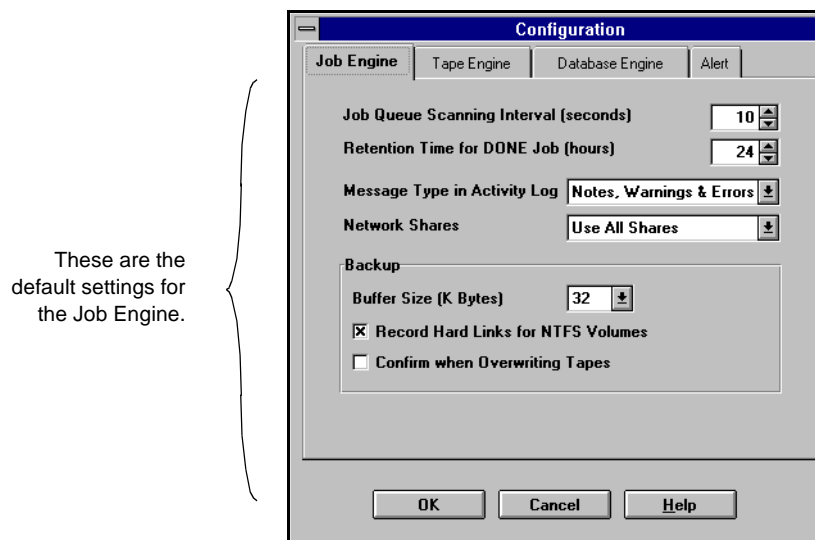
Items that you can configure for the engine appear in this area of the dialog box.



Configuring the ARCserve Job Engine

The ARCserve Job Engine controls when jobs in the job queue should be executed. It scans the job queue regularly, launching jobs as their execution date and time are reached. Follow the directions below to configure the ARCserve Job Engine.

1. In the Configuration dialog box, select the Job Engine tab. The following figure shows the default configuration for the Job Engine:



2. Provide the information to configure the Job Engine and click OK.

Job Engine configuration fields

Job Queue Scanning Interval (seconds) - default is 10 seconds

The Job Engine constantly scans the job queue for jobs that should be executed. By default, the job queue is scanned every 10 seconds. To change the time interval that the job queue is scanned, specify a number between one and 9,999.

Retention Time for DONE Job (hours) - default is 24 hours

Jobs with a final status of “DONE” remain in the job queue for the time that you specify in this field. By default, ARCserve will keep DONE jobs for 24 hours before they are deleted from the queue. To change the time that DONE jobs remain in the queue, specify a number between 0 and 999. For more information on DONE jobs, refer to Chapter 8, “Managing Jobs and the Job Queue”.

Message Type in Activity Log - default is “Notes, Warnings & Errors”

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 that are included in the Activity Log, specify one of the following values:

- None - No messages will appear.
- Errors - Only errors that occur will appear.
- Warnings & Errors - Warnings and errors that occur while running ARCserve will appear.
- Notes, Warnings & Errors - The default selection, includes all notes, warnings, and errors.
- Debug - Only debugging information will appear. This should only be used for troubleshooting purposes.

Network Shares - default is “Use All Shares”

This option applies only for machines on a Windows NT network. The selected option determines how the network is displayed in the Browser. The available options are:

-
- Use All Shares - Display administrative and user shares.
 - Use Default Shares Only - Display only administrative shares.
 - Use Users Shares Only - Display only shares that have been specifically set by users.

For example, if you select Use Default Shares Only, then only default shares are displayed in the Browser. If you back up a machine, you are only able to back up the default share drives for the selected machine.

For more information on shares, refer to your Windows NT documentation.

Backup -> Buffer Size (K Bytes) - default is 32 KB

ARCserve waits for the buffer to fill up with data before it writes to the tape drive. By default, the size of the buffer is 32 kilobytes. To change the size of the buffer used by ARCserve when performing a backup, drop down the list and select one of the available sizes.

Backup -> Record Hard Links for NTFS Volumes - default is on

If you back up files that have hard links, this information is included and preserved by default.

Backup -> Confirm when Overwriting Tapes - default is off

Any time a tape will be overwritten, you can have ARCserve 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 cancelled.

Configuring the ARCserve Tape Engine

The ARCserve Tape Engine identifies the tape drives that are connected to your system and activates them when a job requiring them is started. The steps below describe how you would configure the Tape Engine.

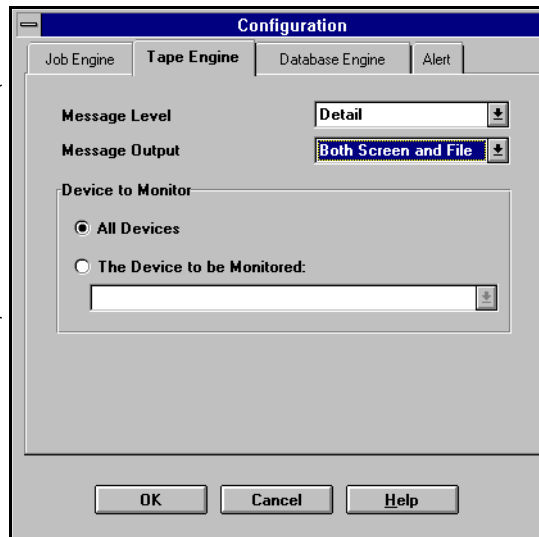


Changing the configuration of the Tape Engine is not recommended unless you are having problems with your tape devices. Keep the default configuration for the Tape Engine, except for troubleshooting purposes.

1. In the Configuration dialog box, select the Tape Engine tab.

The default configuration for the Tape Engine appears as shown below:

These are default settings for the Tape Engine.



2. Provide the information to configure the Tape Engine and click OK.

Tape Engine configuration fields

Message Level - default is "None"

The value you specify for this field determines whether you need to specify values for the other fields in this dialog box. If you keep the default ("None"), you do not need to specify anything else.

If you specify a *Message Level* value other than "None", ARCserve will monitor the Tape Engine and log all information in a Tape Engine Log. The Tape Engine Log will appear in the Quick Status Manager window. If you keep the default "None" for *Message Level*, the Tape Engine Log will not appear in the Quick Status Manager. To change the level of Tape Engine messages, specify one of the following:

- None - The default value, nothing is logged.
- Brief - Non-redundant messages sent by the drives are logged.
- Detail - Every message sent by the drives is logged.

Message Output

If you specified "Brief" or "Detail" in the *Message Level* field, 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 Quick Status Manager.

Device to Monitor

If you specified “Brief” or “Detail” in the *Message Level* field, you can specify which device(s) you want to monitor. To specify the drive to monitor, select one of the following:

- All Devices - All the devices connected to your machine will be monitored.
- The Device to be Monitored: - Select this option to monitor only one device. Drop down the device list and select the one to monitor.

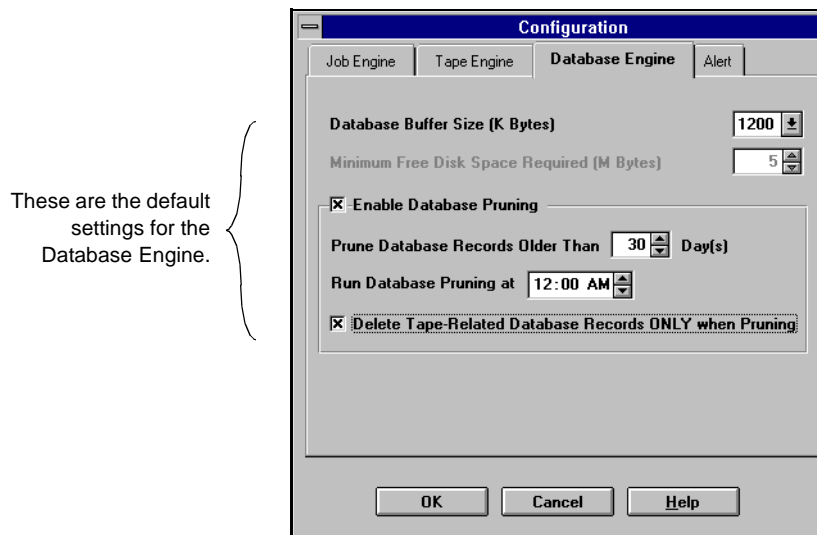
Configuring the ARCserve Database Engine

The ARCserve Database Engine stores all the important information that ARCserve needs. Some of the information saved to the database are:

- what files and directories have been backed up, copied, and restored
- jobs that ARCserve has processed
- tape devices and tapes used for ARCserve operations

1. In the Configuration dialog box, select the Database Engine tab.

The default configuration for the Database Engine is shown below:



2. Provide the information to configure the Database Engine and click OK.

Database Engine configuration fields

Database Buffer Size (K Bytes) - default is 600 KB

By default, ARCserve uses a buffer that is 600 KB for recording information in the database. To change the size of the buffer used by ARCserve when writing to the database, drop down the list and select one of the available sizes.

Minimum Free Disk Space Required (M Bytes) - default is 5 MB

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 for the database. To change the amount of disk space required for the database, specify a number between 1 and 10.

Enable Database Pruning - default is on

Information regarding the files and directories that were backed up in a session is deleted when database pruning is done. By default, this option is selected. It is useful to maintain the detailed information for restoring purposes. However, not pruning your database can mean that it fills up quickly. If you want to free up space in the database file, you might want to set this option on.

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 prunes it. By default, the record will be pruned (if enabled) after it has been in the database for 30 days. To change the length of time a record has to remain in the database before pruning, 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.

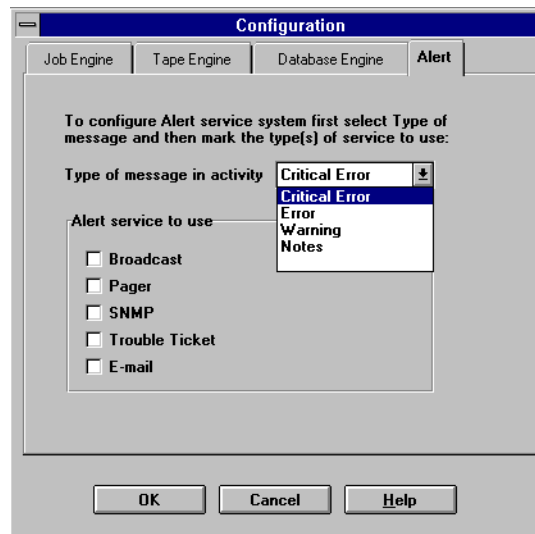
*Delete Tape-
Related
Database
Records ONLY
when Pruning*

When you reformat or erase a tape, ARCserve will also 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.

Alert options

You can use the Alert notification system to send messages about events in your ARCserve Server operations. For more information on the Alert system, refer to Chapter 16, “Using the Alert Manager”.

To display the Alert options, select the Alert tab in the Configuration dialog box.



Event Choose the type of activity for which you want to be notified.

Method Choose one or more of the following Alert notification methods. Before you choose a method, make sure that the Alert system is properly configured for this method. Refer to Chapter 16, “Using the Alert Manager” for more information.

Method	Description
Broadcast	Windows NT Broadcast's can be sent to specific users or groups.

Method	Description
Pager	Alphanumeric paging system.
SNMP	Simple Network Management Protocol (such as NetWare Management System (NMS) and HP Open View).
Trouble Ticket	Any print queue on your network. (Trouble Tickets can only be sent to non-Postscript printers.)
Email	Microsoft mail.

15

Chapter

ABOUT THE ARCSERVE AGENTS

In this chapter, you will learn how the Client Agents can work with ARCServe running on different platforms.

In this chapter, you will learn:

Page:

- 15-2 ➤ What the Client Agents are
- 15-2 ➤ The benefits of using the Client Agents
- 15-3 ➤ How the Client Agents work
- 15-6 ➤ How to enable the Client Agents

What is a Client Agent?

The Client Agent software is an optional product that can be used with your ARCserve for Windows NT software. The Client Agent technology is a fully supported feature of ARCserve for Windows NT. The benefits of utilizing the Client Agent technology along with your ARCserve for Windows NT software are outlined below:

- Using the Client Agent for NetWare software enables all NetWare specific information (such as bindery and trustee information) to be backed up or restored intact.
- Compatibility and support for SPX/IPX and TCP/IP protocols.
- Support for file system "browsing" of the remote clients.
- Maximize your ARCserve server's capabilities and speed by off-loading the processing of the backup, restore, and compare functionality directly to the machine on which the Client Agent resides.



The Client Agents are available only with the ARCserve for Windows NT Enterprise edition. The single server release does not support this capability.

How do the Client Agents work?

15

ARCserve for Windows NT allows you the option of performing complete and safe backup, compare, and restore jobs with or without the Client Agents running. However, Client Agents provide added functionality when you perform these jobs.

Push Agent technology and efficient backups

The Client Agent for NetWare and the Client Agent for Windows NT both utilize the Push Agent technology to accomplish backup, restore, and compare jobs. The advantage of Push Agent technology is that it requires a NetWare server or a Windows for NT machine to perform the necessary filtering of files and directories prior to "pushing" them to the host server (the server running ARCserve) for processing. By utilizing this technology, the transfer of data to the ARCserve host is now minimized thereby reducing load on your communications network. The host server is more efficient because it no longer has to process the files or directories that were flagged for a job. For example:

In your Sales Department, you have a NetWare server that is scheduled for backup. The only directories that you schedule for backup are the user's home directories. If you back up the Sales server without the Client Agent, then everything on that server is sent to the ARCserve host. ARCserve then filters out the specific directories flagged for backup.

Using the Client Agent, the filtering is done directly on the source machine and then only the specified directories or files are sent to the ARCserve host for backup. The ARCserve host's resources are better utilized since the filtering is off-loaded to the source machine.

System Files

An important advantage of the Push Agent technology is that it retains the attributes of system files. This is useful, particularly as it relates to NetWare servers. Because the attributes of the system files are maintained, they are restored as system files rather than as DOS files. This enables you to perform complete disaster recovery.

While a backup done without using the Push Agent technology is as effective and safe, the Push Agent technology allows the restore operation to be completed much more efficiently. Since the system files are backed up intact (native attributes are maintained), they can be restored as is in the event of a disaster scenario, thereby reducing your down time. A restore operation on a server that was backed up without utilizing the Push Agent technology, would require that the system files and bindery be rebuilt. In effect, your NetWare server would be restored but it would be generic. You would still have to re-do the system files and bindery before you were completely operational again.

The advantage of utilizing the Client Agent's use of Push Technology becomes evident. Because the system files and bindery attributes are maintained, the restore operation results in a NetWare server that is immediately ready to use. The system files and bindery do not have to be rebuilt.



The Client Agents can only be used for backup operations for ARCserve for Windows NT v2.0. and above.

Installing and Loading the Client Agents

The client agent must be properly installed and loaded in order for ARCserve for Windows NT to use the agent. See the installation guide for your client agent for more information.

Enabling the Client Agent for ARCserve



Before enabling the Client Agent, you must load it on each machine you will be accessing. Be sure to note the address of the source machine (the machine the Agent is loaded on) when loading the Client Agent. See Chapter 4 of this manual for complete instructions on how to do a backup.

You must enable the client agent before performing a backup, restore, or copy operation.

To enable the agent:

1. Highlight the machine in the Browser of the Backup, Restore, or Copy window.
2. Click the server name with your right mouse button (or choose Source Context Menu from the Backup menu).
3. Choose Use Agent from the popup menu.

The Agent dialog box opens.



4. Click the Use Agent checkbox.
5. Select a protocol and enter the protocol address on the remote machine.
6. Click OK.

The client agent is enabled for this machine.



16

Chapter

USING THE ALERT MANAGER

ARCserve for Windows NT can alert you (or anybody you choose) whenever a virus is detected on your network and when jobs are completed or terminated.

In this chapter, you will learn:

Page

- | | | |
|------|---|---|
| 16-2 | ➤ | What Alert is |
| 16-2 | ➤ | How Alert works |
| 16-4 | ➤ | How to run Alert |
| 16-5 | ➤ | How to configure your system to send alerts |

Alert basics

What is Alert?

Alert is a notification system that sends messages to people in your organization using different communication mechanisms. Alerts can be sent to the system administrator, a hardware technician, or anyone else, in or out of the office.

How does Alert work with ARCserve for Window NT?

Alert does not generate its own messages. 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 network users or groups.
- Pager - Numeric and alphanumeric.
- Electronic Mail - Microsoft Mail must be installed on a server on your network.
- Trouble tickets - An alert can be printed through any print queue on your network. (Trouble tickets can only be sent to non-Postscript printers.)
- Simple Network Management Protocol (SNMP) managers - Such as NetWare Management System (NMS) and HP OpenView.

What are the components of Alert?

Alert has four basic 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.
- *.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 NT directory in order for Alert to handle messages generated by an application.
- Alert Manager - This is where you can configure how Alert broadcasts its messages.

What do you have to do to generate alerts?

In order for ARCserve to generate alerts you must tell Alert what information is necessary to communicate. For example, if you will be using the pager system, you will have to tell Alert what pager number to dial, and you will have to supply information about your modem.

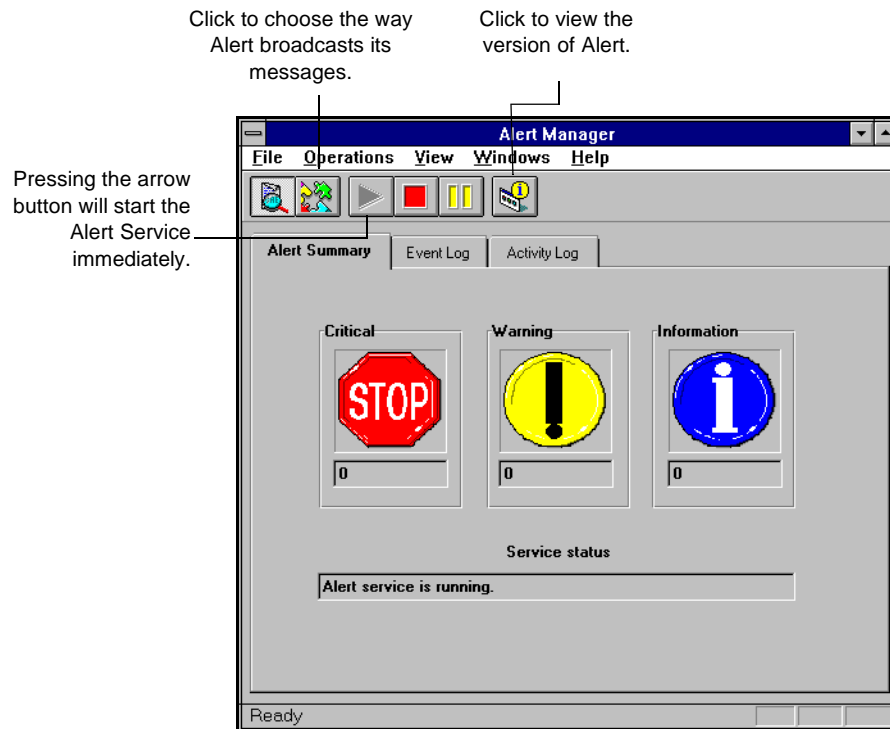
All of this information must be configured in the Alert program on your ARCserve for Windows NT server. This configuration information is shared by all of your domain machines.

Running the Alert Manager

To run Alert:

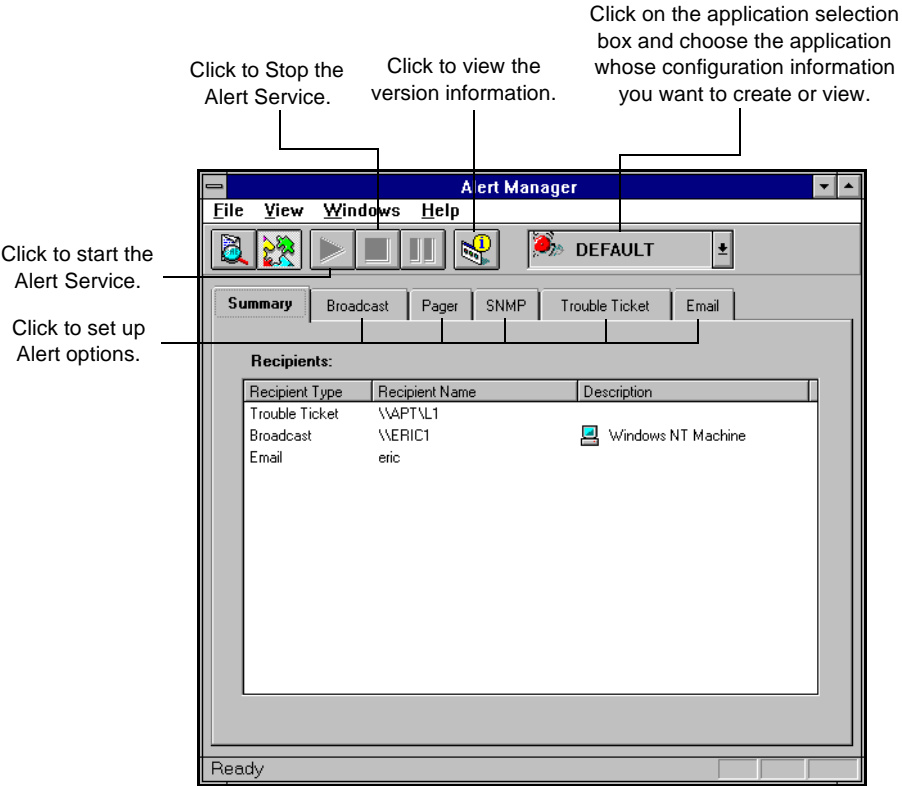
1. Click the Alert Manager icon in the Program Manager.

The Alert manager screen appears:



Configuring Alert

When you enter configuration information in Alert, the information is used for all applications using Alert. You can also enter configuration information specifically for ARCserve. This configuration information will override the default Alert configuration.



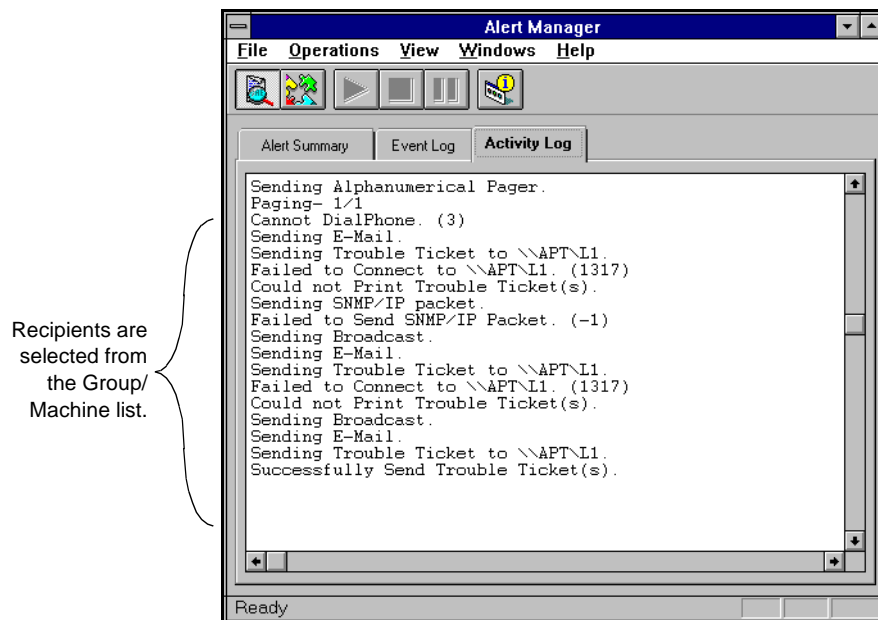
Using the Broadcast option

Alert broadcasts can be sent to specific network users or groups.

Adding broadcast recipients

To add broadcast recipients:

1. In the Alert Manager, click the Broadcast tab.
A list of all groups and recipients appears:

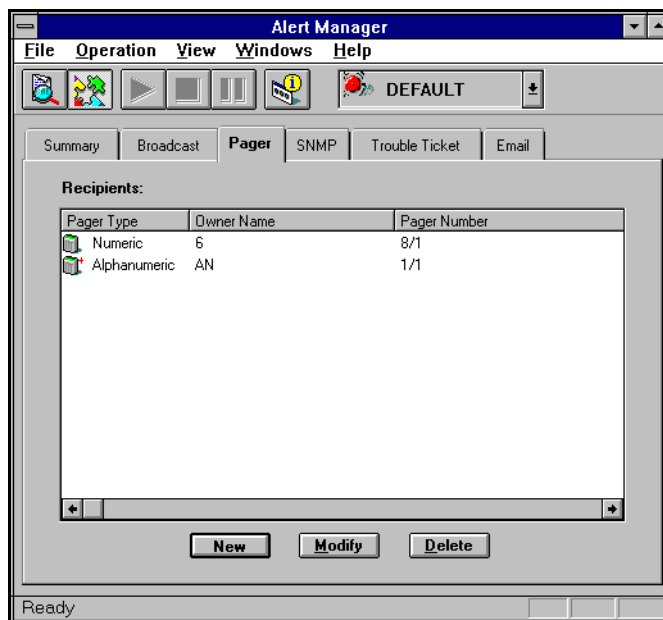


2. Highlight the group/machine you want.
3. Press the Add button.

Using the Pager option

The pager option is used to send a numeric or alphanumeric pager message.

When you click the Pager tab, the current list of recipients appears:



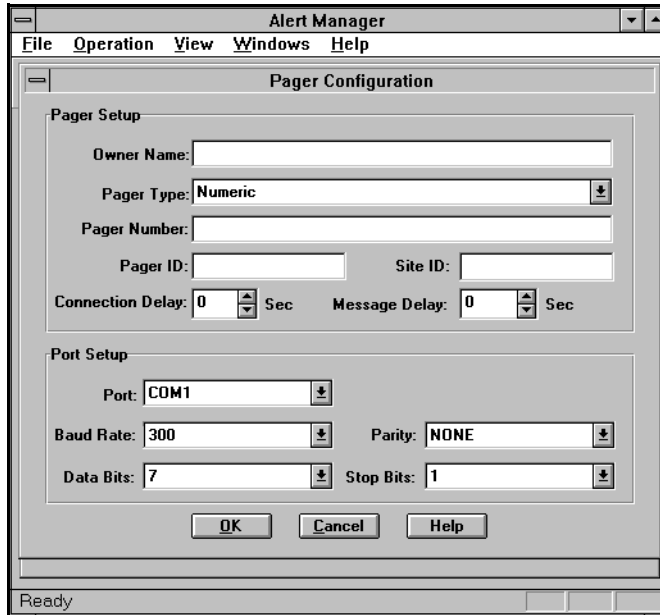
The list is initially empty.

Adding pager recipients

To add pager recipients:

1. Press New.

The Pager Configuration screen appears:



2. Enter information on the Pager Configuration screen .

Owner name

Enter the name of the pager recipient.

Pager Type

Indicate if you are using a numeric or alphanumeric pager.

Pager Number

Enter a maximum of 24 characters. If a digit, such as 9, must be dialed to get a dial tone, it must be included in this field.

A comma can be entered to indicate a one second pause. If a longer pause is desired, a string of commas can be entered.

A dash (-) can be used to separate digits, but it has no function. (Since this can vary by modem, you should verify this with your modem manual.)

Pager ID

Enter up to eight digits to identify the pager that will receive the alerts.

<i>Site ID</i>	Enter up to four digits to identify where the alert occurred. This number is included in the message to the pager. Therefore, if the number is less than four digits, you should use leading zeros.
<i>Connection Delay</i>	Enter the number of seconds for each delay you want to wait before a connection is made with the pager company. When a phone number is dialed, a connection is made. The connection may be established immediately or it may be delayed. This will vary with your pager company, location, time of day, telephone equipment, and telephone traffic. If the connection is not established immediately, adding a delay can prevent the alert from being sent before the connection is established.
<i>Message Delay</i>	Enter the number of seconds for each delay you want to wait between the time the connection is made and the alert message is sent.
<i>Port</i>	Enter the communications port you want the pager message to be broadcast from.
<i>Baud Rate</i>	Indicate the baud rate being used by your modem.
<i>Parity</i>	Indicate the parity setting, none, odd, or even, of your modem.
<i>Data Bits</i>	Enter the number of data bits, 7 or 8, that your modem uses.
<i>Stop Bits</i>	Enter the number of stop bits, 1 or 2, that your modem uses. * For alphanumeric pagers, the recommended usage is 7 data bits, even parity, 1 stop bit; or 8 data bits, No parity, 1 stop bit.



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

3. Press OK to save your information .

Interpreting the pager message

Numeric pagers

When a numeric pager is sent a virus alert, the coded message will appear as: **Message = DDSSSSCC**

DD is the virus detection code number. It tells you which component of ARCserve has detected a virus.



You must check ARCserve's scanning records to determine which machine is infected and which files or directories contain the virus.

Virus Detection Code	Description
01	N/A
02	N/A
03	A boot virus or a change to the Critical Disk Area was detected on a local machine.
04	The ARCserve Manager detected a virus at a local machine.
05	The ARCserve Server detected a virus on a machine.
07	N/A
08	IMMUNE detected a boot virus.
10	N/A
11	N/A

SSSS is the user defined site number from the Pager Configuration screen. The site number represents the machine that detected the virus.

CC is the user defined custom code from the Pager Configuration screen. The custom code represents the machine that sent the message.

Alphanumeric
pagers

There are several messages that can be sent to an alphanumeric pager. Each message is listed below. Words that appear in italics are variables that 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*

* This can be the path on a server or workstation.

Using the SNMP option

The SNMP option is used to send an SNMP ‘trap’ (message) to an SNMP manager.

Examples of SNMP managers include NetWare Management System (NMS) and HP OpenView.

This section explains the SNMP fields on the SNMP Configuration screen.

When you click the SNMP tab, the current list of managers appears:

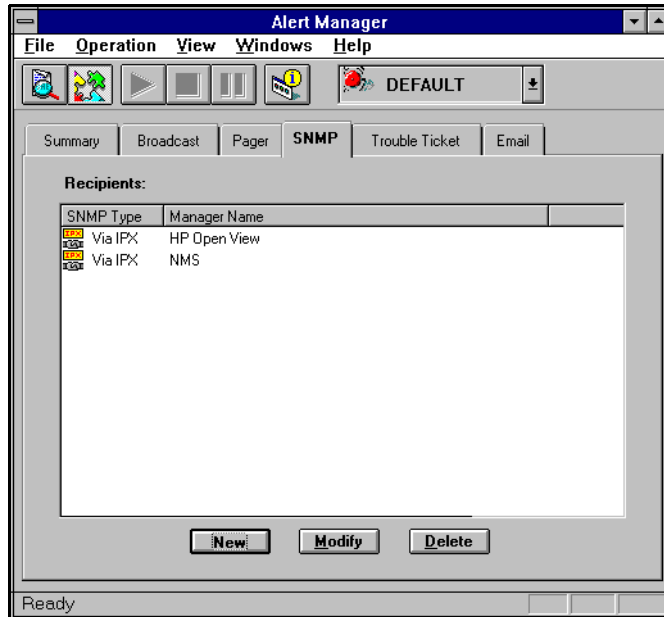


Adding SNMP recipients

To add SNMP recipients:

1. Click the New button.

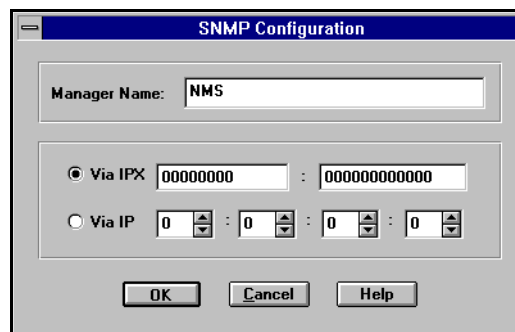
The SNMP Configuration screen appears:



2. Enter information on the SNMP Configuration screen.

Manager Name

Enter the name of the SNMP Manager.



Via IPX

Enter the 8 byte network address of the machine where the SNMP manager is located. Next, enter the 12 byte node

address of the machine where the SNMP manager is located. Use this field for Novell networks.

Via IP

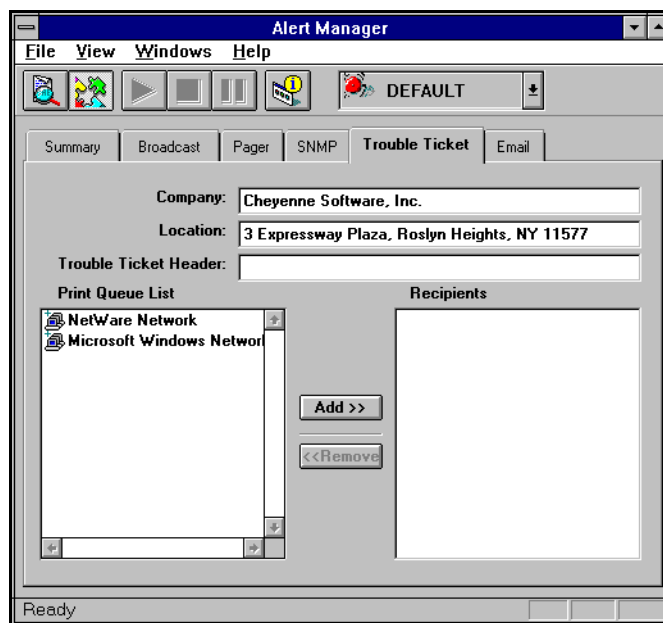
Enter the IP address of the machine where the SNMP manager is located. Use this field if you are running the TCP/IP stack.

3. Press OK to save your information .

Using the TroubleTicket option

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

1. Click the Trouble Ticket tab.



Company Name

Enter your company name.

Location

Enter the location of your company.

Trouble Ticket Header

Enter the information that should appear at the top of each Trouble Ticket.

Choosing the Printer

1. Choose a printer from the Print Queue List by highlighting it.

2. Click the Add button to move the printer to the Recipients column.

Using the E-mail option

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



Microsoft Mail 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 manual for instructions about how to set up your E-mail account.

Click this box if you want Alert to notify you when the E-mail is received.

Click this to attach files (.doc, .bmp, etc.) to your E-mail document.

Click this to remove files attached to your E-mail.

Return receipt

Address

To:

CC:

Subject:

Attach

Detach

File Name:

Ready

Adding E-mail recipients

There are two steps to add E-mail recipients:

1. Log into E-mail.
2. Select E-mail recipients.

Log into E-mail

1. Click the Address button.
You will be prompted for your Microsoft Mail password.
2. Enter the user name and password Alert should use when logging into Microsoft Mail.

Using your Alert PostOffice to select E-mail recipients

Once you have logged into E-Mail, the Alert PostOffice screen will appear. This screen contains the names of users in your post office.

Click to select a user directory from the list of available directories for PostOffices and gateways on your network.

Click to display names and groups in your personal address book.

Click to find a name in the current user directory.

Click to add a new name and address to your personal address book or to address a message to another E-Mail system. You can add a gateway address plus the name and address of a user on another E-Mail post office.

The screenshot shows the 'ALERT PostOffice' window. At the top, it says 'Directory: Postoffice List'. Below this is a list of names: Andrew Goldberg, Brian Gaughran, Chris Hart, Diane Lang, Jon Goddard, Karla Zamborino, Mark Rapp, Megan Elizabeth Jerlin (highlighted), Nathalie Trouillot, and Pramod Mari. To the left of the list are icons for address book, search, and a speech bubble. Below the list are buttons for 'Add:', 'To', 'Cc', and 'Details'. At the bottom are 'To:' and 'Cc:' fields, and 'OK', 'Cancel', and 'Help' buttons.

Click to see detailed information about the highlighted user.

1. Locate usernames in the Alert PostOffice.
To locate usernames in the Alert PostOffice you can:
 - Enter the first few characters of a name. The list scrolls to the first name in the list with those characters.
 - Use the scroll bar or press the arrow keys to scroll through the list.

2. Select a username.

Highlight the name you have selected and click the To button. The selected name will appear in the To box.

To add names to the Cc list, repeat the steps for locating and selecting names, and then click the Cc button.

3. Click OK.

Alert's Activity Log

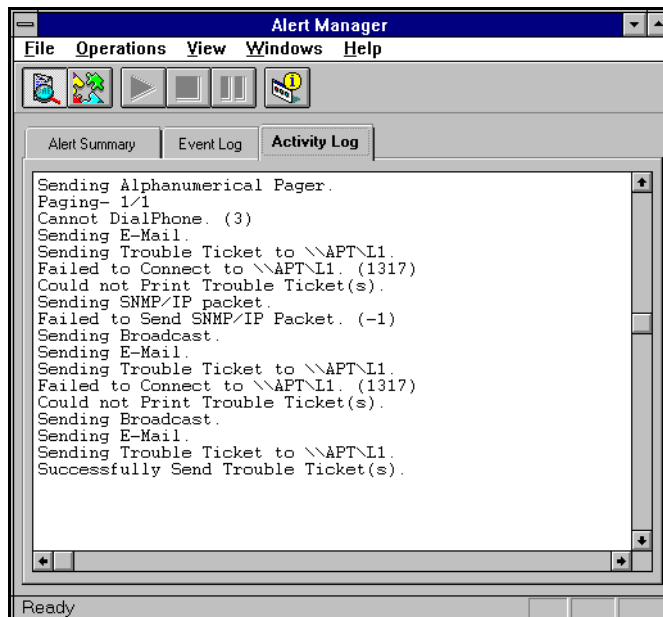
While the current status of Alert is displayed on Alert's main menu, a historical listing is stored in the Activity Log. You can view, print, or clear this log.

Displaying the Activity Log

To display the Activity Log:

1. Select the Activity Log Tab from the Alert manager.

Your Activity Log, similar to the one shown below, appears:



Clearing the Activity Log

You can delete the Activity Log. You might want to do this if Alert has been running for a long time and the log has grown large. To clear the log:

1. Select Clear Activity Log from the Operations menu.

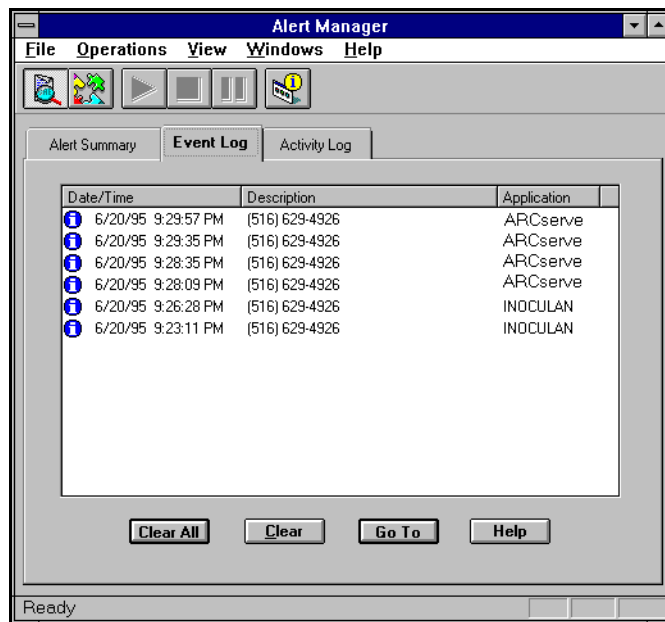
Alert's Event Log

Every message that is generated by Alert is stored in the Event Log. You can view, print, or clear this log.

Displaying the Event Log

To display the Event Log:

1. Select the Event Log Tab from the Alert Manager.



Clearing the Messages Report Log

You can delete the Event Log. You might want to do this if Alert has been running for a long time and the log has grown large. To clear the log:

1. Click the Clear All button.

If you want to clear one item, highlight it and click the Clear button.

Fields on the Event Log screen

<hr/> <i>Date/Time</i>	Displays the date and time the event occurred.
<hr/> <i>Description</i>	The description is determined by the applications that set the alert.
<hr/> <i>Application</i>	Tells you the application that generated the event.

