

User's Guide and Technical Reference

TapeWare®

Yosemite Technologies, Inc

Information in this document is subject to change without notice. Yosemite Technologies, Inc makes no representations or warranties with respect to the contents of this document and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Further, TapeWare reserves the right to revise this publication and to make changes without obligation to notify any person or organization of such revision or changes.

The persons mentioned in the examples of this manuscript are historical figures associated with Yosemite National Park. Beyond the use of their names, the examples in this text are fictitious. Any other resemblance, incidental or otherwise, to any person living or dead are purely coincidental.

TapeWare is a registered trademark of Yosemite Technologies, Inc.

NetWare is a registered trademark of Novell, Inc.

Windows and Windows NT are registered trademarks of Microsoft Corporation.

Windows Explorer is a trademark of Microsoft Corporation.

Macintosh is a registered trademark of Apple, Inc.

IBM and OS/2 are registered trademarks of International Business Machines Corp.

© 1998-2000 Yosemite Technologies, Inc. All rights reserved.

Contents

Before You Begin xix

Before You Install TapeWare xix
TapeWare Documentation xix
On-line Help xx
Customer Support xx

Overview of TapeWare Concepts 1

Physical Peripherals 1

The Storage Management Database 2

The Backup Management Program 3

Managing the Storage Management Database 3
Objects 3
Properties 4
Storage Management Zones 5
Creating and Running Jobs 6

The TapeWare Workplace 9

In This Chapter 9

Installing TapeWare 9

Product Configuration 10
Installation and Storage Management Zones 10
Installation Instructions 11
DOS 11
NetWare 11
Windows (95/98 or NT) 11
The TapeWare Installation Manager 12
Install TapeWare 12
Remove TapeWare 12
Install an Option 12
Remove an Option 12
Update License 12
Installing TapeWare as a Service 13
Installing Backup Devices 13

Starting TapeWare 14

The Logon Window 14

Selecting a Storage Management Zone 14
User Name and Password 15
Changing Your Password 16
Logging on the First Time 17

<i>Grace Logons</i>	17
Logging Out	17
<i>Logging Out and Running Scheduled Jobs</i>	18
The Main TapeWare Object Window	18
<i>Tabs</i>	19
<i>Toolbar</i>	19
<i>The Object View Areas</i>	20
<i>The Status Bar</i>	21
Working with Objects in the Main TapeWare Window	21
<i>Arranging Objects</i>	22
<i>Keyboard Shortcuts</i>	22
<i>Shortcut Menus</i>	23
The Find Button	23
Property Sheets	23
Opening Property Sheets	24
Menus	25
Reports	26
Additional Report Capabilities	28
Working with the TapeWare Wizard	28

Creating TapeWare Jobs 31

<i>In This Chapter</i>	31
Overview	31
Backup, Restore, and Verify Jobs	31
Backup Jobs	32
Restore Jobs	32
Verify Jobs	33
The Job and Database Tabs	33
Creating New Jobs	33
Creating Jobs with the Wizard	34
Creating New Jobs from a Job Tab	34
Creating New Jobs by Copying	36
Renaming, Deleting, and Moving Jobs	37
Organizing Jobs with Folders	37
<i>Types of Folders</i>	38
<i>Folders and Job Tabs</i>	38
The Home, Admin, and Everyone Folders	39
<i>The Everyone Folder and Permissions</i>	40
Creating Job Folders	40
Moving, Renaming, and Deleting Folders	41
<i>Modifying Folders</i>	41

Permissions 43

<i>In This Chapter</i>	43
------------------------	----

Overview 43**Users and Groups 44**

The Everyone Group 44

Effective Permissions 44

Determining Effective Permissions 45

Examples of Determining Effective Permissions 45

Effective Permissions Example #1 45

Effective Permissions Example #2 46

How Moving Objects Effects Permissions 47

Types of Permission 48

Supervisor 48

Access 49

Create 49

Modify 50

Delete, Read, and Write 50

Examples of Permissions 50

Example #1 51

Example #2 52

Granting Permissions to Other TapeWare Users 53

Restricting User Permissions 54

Selecting Files and Instances 57

In This Chapter 57

Overview 57**Backup Selection Concepts 58**

Marking Files for Backup 58

Selecting Folders vs. Selecting Files 59

Sorting Files with Filters 60

Filters Exclude, not Include 60

Filters and Grayed Folders 61

Selecting Changed Files only 61

Automatically Selecting New Files for Backup 61

Selecting Files for Backup Jobs 62

Applying Filter Criteria 62

The Filter Selection Criteria 63

Backup Range 64

Modify Range 64

Create Range 65

Access Range 65

Size Range 65

Instance Range 65

Must Match 66

Cannot Match 66

Wildcard Type 66

Required Attributes 66

Exclude Attributes 67

Parents 67

Children 67

Media 67

Restore Selection Concepts 68

Selecting Files for Restoring 68

Selecting Instances of Files 69

Instances and Filters 70

Selecting Instances and Folders 70

Selecting Folders vs. Selecting Files 72

Selecting Files and Instances with Filters 72

Filters Exclude, not Include 73

Changing the Name and Location of Restored Files 73

Selecting Instances of Files for Restore Jobs 73

Selecting Specific Instances 74

Applying Filter Criteria 74

The Filter Selection Criteria 75

Backup Range 75

Modify Range 76

Create Range 76

Delete Range 76

Access Range 78

Size Range 78

Instance Range 78

Must Match 78

Cannot Match 79

Wildcard Type 79

Required Attributes 79

Exclude Attributes 79

Parents 79

Children 80

Media 80

Restoring Files with New Names and Locations 80

Restoring a File with a New Name 81

Restoring Files to a Different Folder 81

Restoring Files to a New Folder 82

Verify Selection Concepts 82

Selecting Files for Verifying 83

Selecting Instances of Files 84

Selecting Folders vs. Selecting Files 84

Selecting Files and Instances with Filters 85

Filters Exclude, not Include 85

Selecting Files for Verify Jobs 85

Selecting Specific Instances 86

Applying Filter Criteria	86
The Filter Selection Criteria	86
<i>Backup Range</i>	87
<i>Modify Range</i>	87
<i>Create Range</i>	87
<i>Access Range</i>	87
<i>Size Range</i>	87
<i>Instance Range</i>	87
<i>Must Match</i>	88
<i>Cannot Match</i>	88
<i>Wildcard Type</i>	88
<i>Required Attributes</i>	88
<i>Exclude Attributes</i>	88
<i>Parents</i>	89
<i>Children</i>	89
<i>Media</i>	89

Scheduling Jobs 91

In This Chapter 91

Overview 91

Backup Job Scheduling Concepts 91

Rotating Media 92

Backup Mode 93

Backup Job Schedule Type 93

Selecting a Backup Job Schedule 95

Not Scheduled Backup Jobs 96

Built-In Schedules for Backup Jobs 96

Which Built-in Schedule to Select 97

Full Data Recovery Period 97

Access to Historical Files 97

Minimum Number of Tapes or Media 98

Built-In Schedules Compared 98

Scheduling Backup Jobs with Built-in Schedules 99

Customizing Schedules 100

Creating Custom Schedules 100

Modifying Built-in Schedules 101

Modifying the Custom Schedule 102

Changing the Tape Count 102

Changing the Backup Mode 103

Determining the Length of the Full Data Recovery Period 103

Incremental Jobs and Full Data Recovery 104

Manual Backup Jobs 104

Scheduling Restore Jobs 105

Scheduling Verify Jobs 106**Job Options 109***In This Chapter 109***Overview 109****The Backup Job Options Tab 109***Backup Mode 110**Write Mode 111**New Media Name 112**Auto Format Mode 112**New Media Location 113**Number of Passes 114**Delay Between Passes 115**Log option 115**Auto Verify Mode 116**Software Compression 117**Device 117**Media 118**Media Password... 118**Change Mode 119**Backup Options Automatically Updated 120***Advanced Options 121***Native data streams format 121**Parent Security 122**Child Security 122**Volume restrictions 122**Space restrictions 122**OS/2 extended attributes 123**Macintosh Finder 123**NFS information 123**Object owner 123***Restore Job Options 123***Number of Passes 124**Delay Between Passes 125**Log options 125**Device 125**Advanced Options 126***Verify Job Options 127***Number of Passes 127**Delay Between Passes 128**Log options 128**Verify Mode 128**Device 129*

Running Jobs 131

In This Chapter 131

Overview 131

The Queue Tab 132

The Queue Tab Detail View 132

Running Scheduled Jobs 132

Automatically Running Scheduled Jobs 133

Security and Scheduled Jobs 134

Forcing Scheduled Jobs to Run 134

How Forcing Jobs to Run Effects Job Parameters 135

How Forcing Jobs to Run Effects Permissions 136

Viewing and Printing Scheduled Job Instructions 137

Running Unscheduled Jobs 138

Unscheduled Job Parameters and Permissions 138

The Job Status Window 139

Job Status Messages 140

Viewing and Printing the Job Log 141

Email Support for Job Logs 143

Viewing and Printing Audit Logs 144

The Media, Device, and Database Tabs 147

In This Chapter 147

Overview 147

The Media Tab 148

Creating Media Folders 148

Deleting Media Folders 149

Creating New Media 149

Media Passwords 151

Deleting Media 152

The Device Tab 152

Restarting Failed Devices 153

Device Tab Commands 153

Eject Media 154

Eject Magazine 154

Rewind 154

Retention 154

Clean Device 154

Erase Quick & Secure Erase 155

Format Media 155

Identify Media 157

Import Media 157

Restore Database 158

The Database Tab 161

Tips, Techniques, and Strategies 163

In This Chapter 163

Managing the Storage Management Database 163

- Where to Locate the Storage Management Database 163
- Accounting for Storage Management Database Size 166
- Storage Management Server and Machine Platforms 167

Strategies for Faster Jobs 167

- What Slows Down Tape Drives? 167
- Maintaining the Flow of Data 168
- Other Factors that Effect Job Speed 170

Working with Permissions 170

- Checking the Effective Permissions of a User 171
- Using Groups to Handle Complex Security Needs 171

Working with Scheduled Jobs 172

- Running a Failed Rotation Job Again 172
- Scheduling a Job to Run Once Only 173

Selecting Files for Jobs 174

- Selecting Files not Previously Backed Up 174
- Selecting Deleted Files for Restoring 175
- Selecting Instances from a Specific Job 175
- Selecting Instances from Specific Media 176

Restoring Tips 176

- Restoring Volumes for the Latest Date 176
- Restoring Volumes for a Specific Date 176
- Copying a Directory Structure 178
- Restoring Files to a New or Different Folder 179
- Restoring Files with New Names 180

Other Tips 181

- Moving Data between Operating Systems 181
- Setting up an Autoloader for Automatic Cleaning 182

Permissions and Security Reference 183

In This Chapter 183

Overview 183

- Before Proceeding Further 185

Adding New Users and Groups 185

- New User/Group Folders 186
- Setting Up Users 186

Logon Control Tab, User Object 186

Groups Tab, User Object 188

Equivalencies 188

Permissions Tab, User Object 189

- Setting Up Groups 191

Members Tab 191

<i>Permissions Tab, Group Object</i>	191
Effective Permissions	192
Calculating Effective Permissions	192
<i>Effective Permissions Algorithm</i>	193
<i>Permissions from Multiple Sources</i>	193
Examples of Effective Permissions	194
Checking Effective Permissions	196
Permissions Reference	197
Read Permission	198
Write Permission	198
Delete Permission	199
Modify Permission	199
Create Permission	200
Access Permission	201
Supervisor Permission	201

Objects and Properties Reference 203

Address Tab	204
Protocol	204
Net Address	204
Audit Tab	205
Command Tab	206
Command to execute before backup	206
Command to execute after successful backup	206
Command to execute after failed backup	207
Communication Test Tab	208
Transfer type	208
Communication Status	208
With data checking	209
Connections Tab	210
Active Connections	210
Details	210
Diagnostics Tab	211
Drivers Tab	212
Email Tab	213
Equivalencies Tab	214
Users to which this user is equivalent	214
Users to which this user is not equivalent	215
General Tab	216
Name	216
Type	216
Effective Permissions	217
Attributes	218
Size	218

Created	218
Deleted	219
Modified	219
Accessed	219
Groups Tab 220	
Groups to which this user belongs	220
Groups to which this user does not belong	220
Instance Window 222	
Available instances	222
Details	223
Loader Test Tab 224	
Status	224
Total Moves	224
Logon Control Tab 225	
Expiration	225
Grace logons	226
Password	226
Connection	226
Logs Tab 228	
Logs for this object	228
Media Control Tab 229	
First format	229
Last format	229
Last read	229
Mbytes read	229
Count read	230
Last write	230
Mbytes written	230
Count written	230
Number of sessions	230
Current Kbytes	230
Media identifier	230
Media Volume tag	230
Members Tab 231	
Members who belong to this group	231
Members who do not belong to this group	232
Options Tab 233	
Number of passes	233
Delay between passes	233
Log options	234
Backup mode	234
Change mode	235
Write mode	235
Auto verify mode	236
<i>Software Compression</i>	236

- Auto format mode 237
- New media location 237
- New media name 238
- Device 238
- Media 239
- Media Password... 239

Permissions Tab 241

- Users or groups which have permissions to this object 241
- Objects to which this user or group has permissions 242
- Permissions 242

Ping Test Tab 243**Preferences Window 244**

- Confirmation 244
- Tree View 244
- Text editor 245

Schedule Tab 246

- Type 246
- Start time 247
- Daily, Weekly, Monthly, Yearly 247
- End of Week 247

Selection Tab 248**Selection Filters Window 249**

- Backup range 250
- Modify range 250
- Create range 250
- Delete range 250
- Access range 250
- Size range 250
- Instance range 250
- Wildcard type 251
- Must match 251
- Cannot match 251
- Required attributes 251
- Exclude attributes 251
- Parents 252
- Children 252
- Media 252

Status Tab 253

- Loader statistics 253
- Import... 254
- Format... 254
- Identify 254
- Rebuild 254
- Clean 254
- Change Status... 254

Storage Tab 255

Backup stream 255

Audit object actions 256

Query Window 257

Backup range 257

Modify range 258

Create range 258

Delete range 258

Access range 258

Size range 258

Instance range 258

Wildcard type 258

Must match 259

Cannot match 259

Required attributes 259

Exclude attributes 259

Parents 259

Children 260

Media 260

Configuring NetWare Btrieve Database Sets 261*Defining Btrieve Sets 261**Backing up and Restoring Files 263**Additional Notes 263***Configuring Email Support 265**

Installing Email 265

Configuring Email Packages 265

*The MAPI Email Configuration Window 266**The SMTP Email Configuration Window 266*

The Email Tab 267

Working with Microsoft Exchange Server databases 269*In This Appendix 269***Microsoft Exchange Configuration Window 270***Transfer buffer size 270**Force modes 270***Microsoft Exchange Server Notes 271**

Microsoft Exchange and Windows NT 272

Backup Modes 272

*Backup Modes and Circular Logging 272***Restoring Microsoft Exchange Databases 273**

Working with Microsoft SQL Server Databases 275

In This Appendix 275

Overview 275

SQL Configuration Window 276

Admin name 276

Admin password 276

Log level 276

Xfer buffer size 277

Force modes 277

SQL Server Backup Job Notes 278

SQL Server Databases and the Backup Mode 278

Backup Modes 278

Additional Conditions 278

Using TapeWare with SQL Server's Backup Routine 279

SQL Server Restore Job Notes 279

Restoring SQL Server Transaction Logs 279

Restoring SQL Databases to Devices 280

Restoring SQL Databases with a New Name 281

Restoring SQL Server User Databases 281

Restoring SQL Server Master Databases 283

Step 1 - Rebuild the master Database 285

Step 2 - Restart SQL Server in Single-user Mode 286

Step 3 - Restore the Master Database from the Most Recent Backup 287

Step 4 - Apply Changes to the Master Database 287

Step 5 - Drop Invalid Databases And Database Devices 288

Step 6 - Restore the msdb Database 288

Troubleshooting Guide 291

In This Appendix 291

Troubleshooting Installation 292

1. I entered my Key Code correctly, but the installation won't continue. 292

2. I selected 'Connect to an existing storage management server' during installation, but now that storage management server is not displayed. 292

3. When attempting to log on to NDS, TapeWare keeps returning to the password prompt. 292

4. I can't log on to TapeWare. 293

Troubleshooting Backup Jobs 293

5. When I run a backup job, the backup device used is not the one on the server, but rather the one on my local machine. 293

6. Does TapeWare support media rotation? 293

- 7. TapeWare is running, but I don't see any network mapped drives. 293*
- 8. My scheduled job is not running. 294*
- 9. I have two tape drives, but only one is used to run a job. 294*
- 10. How do I replace media in a rotation group? 294*
- 11. How can I tell when the next job will run and which media is required? 295*
- 12. How can I easily view the logs for each Job? 295*
- 13. How can I determine which files were not backed up? 295*
- 14. Can I print the reports or logs under NetWare? 295*

Troubleshooting Restore Jobs 295

- 15. I can't restore a backup made from Windows NT to Windows 95/98. 295*
- 16. How can I restore data to a different file name? 296*
- 17. How can I restore data to another location? 296*
- 18. How can I restore all the files from a single session? 296*
- 19. How can I determine which files are on particular media? 296*
- 20. When restoring, I get many alerts. What is wrong? 296*
- 21. Does TapeWare back up files as compressed? 297*
- 22. Can I restore NT files to NetWare? 297*
- 23. Can I restore NetWare files to NT? 297*

Troubleshooting Verifying 297

- 24. I occasionally get a 'stream sync error' when verifying media. 297*

Troubleshooting Backup Devices 298

- 25. I can't see my tape drive on the Device tab. 298*
- 26. My autoloader under Windows NT is not recognized by TapeWare. The device is displayed, but no loader. 298*

Troubleshooting the Storage Management Database 299

- 27. How do I choose my storage management database location? 299*
- 28. When recovering from a storage server failure, do I get all the information from the TapeWare database back? 299*

Troubleshooting Error Messages 299

- 29. When attempting to add a new object or during a backup, an 'error 51-Database corrupt' is displayed. 299*

30. An 'error 212–Unknown error' is returned when I restore files under Windows NT. 300

Disaster Recovery 301

- Requirements 301
- Installing Disaster Recovery 302
- Using HP OBDR 302
- Configuring Disaster Recovery 302
 - Windows - Diskettes 303*
 - Windows – CDROM 303*
 - NetWare - Diskettes 304*
 - NetWare – CDROM 304*
- When to Re-create the Boot Diskettes 305
- Preparing for Disaster 305
- Recovering from a Disaster 306
- OBDR Compatibility 309
- Using Disaster Recovery with Autoloaders 309
- Troubleshooting – Preparing for Disaster 310
- Troubleshooting – Recovering from a Disaster 311

Index 313

Before You Begin

The *User's Guide and Technical Reference* provides all of the information necessary to effectively use and implement all of TapeWare's advanced features.

Before You Install TapeWare

Installing TapeWare is simple and easy to do. Prior to installing TapeWare, however, you must make some important decisions about implementing TapeWare. Before you install TapeWare, be certain you know the following:

- the name of storage management zone you wish to create.
- the machine that will be the storage management server.
- which workstations or file servers will belong to the storage management zone.
- the Key Code for the product and configuration you have purchased.
- whether you will install TapeWare as a service on Windows 95/98 or NT.

Detailed instructions on how to install TapeWare can be found in "Installing TapeWare," Chapter 2, including a discussion of installing TapeWare as a service. If you are uncertain about creating a database zone or assigning machines to a storage management zone, read Chapter 1. Two sections in Chapter 10, "Managing the Storage Management Database" and "Strategies for Faster Jobs" provide additional information about locating the storage management database and assigning machines to storage management zones that you will want to consider before installing TapeWare.

TapeWare Documentation

On the Installation CD-ROM, you will find the file Usersgd.pdf along with the Adobe Acrobat Reader 3.0. This file contains the exactly the same contents as this manual, the *User's Guide and Technical Reference*.

If you require extra copies of this manual, you can print the PDF file. It was designed to be printed one-sided, with a margin for placing it in a binder. For consistency with this manual, the page numbers, page layout, table of contents,

and index are the same (except that the PDF file does not have facing pages or even/odd headers).

To print the documentation, open the Usersgd.pdf file with the Adobe Acrobat Reader. Choose **Page Setup...** from the **File** menu and set the proper options for your printer. Be certain to select the **Larger Print Area** option if your printer supports this option. Then choose **Print...** from the **File** menu and print the document.

The contents of this manual are also available in the on-line help file.

(You can also purchase additional printed and bound copies of this manual. Contact us at the numbers listed below for additional information.)

Additionally, you will find a PDF file named Errcodes.pdf. This document, over 40 pages long, contains a complete listing of all the TapeWare error codes. You can use the Adobe Acrobat Reader to print this file.

On-line Help

To get on-line help while using TapeWare, select **Help Topics** from the **Help** menu.

For context-sensitive help while using TapeWare, either

- click on the **Help** button and then click on the part of the object you want help with, or
- click on an object with the right mouse button and select **What's this?**, or
- use the **Tab** key to "target" an area in the TapeWare window and press **F1**.



The Help
button

Customer Support

You can get customer support for TapeWare in one of four ways:

- Visit our web-site at www.TapeWare.Com. (For a quick connection, select **Web Page** from the **Help** menu.)
- Email us at Support@TapeWare.Com.
- Fax us at (559) 292-8908
- Phone us at (559) 292-8888.

Overview of TapeWare Concepts

TapeWare provides a powerful, yet cost effective and easy to use, management tool for protecting data on network file servers and workstations. TapeWare provides users with the capacity to back up and restore data across a network and the ability to administer a comprehensive backup plan.


A complete network backup system consists of three parts: *the physical peripherals, the storage management database, and the backup management program.*

Physical Peripherals

TapeWare works with your existing computer network. A network connects workstations and file servers together in order to allow various users to work together on projects and with common files. Networks also allow users to share peripherals, such as disk drives, printers, fax machines, and modems. Sharing peripherals across a network makes economic sense, since multiple workstations can use a single peripheral. Backup devices, such as tape drives, which backup or copy files onto tapes or other media, may also be shared.

For further information about backup device LAN location, see "Strategies for Faster Jobs," Chapter 10.

Sharing a backup device or tape drive makes operational as well as economic sense. Rather than back up each workstation individually, a shared backup device on a network can back up every workstation and file server on the network. This lowers costs and makes it possible to centralize backup operations for the network. Further, a single user, such as the network administrator, can have primary responsibility for backing up all the file servers and workstations on a network.



Although sharing a backup device or tape drive on a network makes economic and practical sense, it also poses several problems.

- First, security is an issue. Most networks have elaborate security systems that prevent sensitive or confidential data from being accessed by unauthorized users. However, unless protective measures are taken, once these files are backed up onto a tape or other media, any user in physical possession of the media can gain access to those files. Although it is possible to physically store the media in a secure location, a complete network backup system will prevent unauthorized viewers from accessing confidential or classified information.
- Second, tracking the location of backed up files is a further issue. While a single user may be able to find a file backed up onto a floppy disk by manually searching through a stack of disks, this approach is unworkable for large networks. Without the appropriate software, locating a given occurrence of a file may be impossible, since there may be hundreds of thousands of files backed up on hundreds of tapes created over weeks or months.

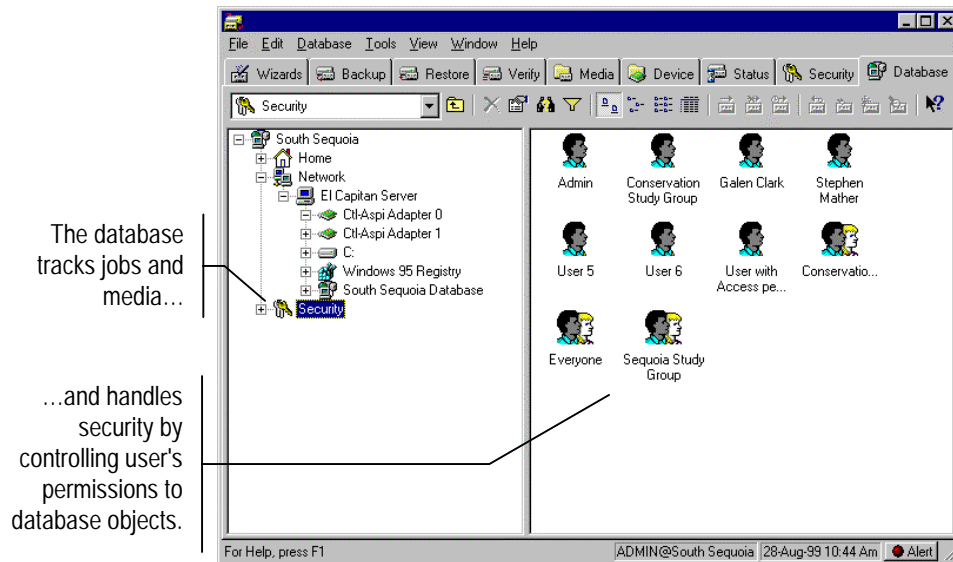
The Storage Management Database

For further information on working with the storage management database, see "Managing the Storage Management Database," Chapter 10.

To solve the problems of security and tracking files, TapeWare creates and manages a **storage management database**. The TapeWare database keeps track of each tape or other media and each file that is on that tape. For each tape, the storage management database contains detailed information about the tape, such as when it was created and who created it, and about the files on the tapes, such as when they were backed up and on which tape the file is stored. This is true as well for other media.

The storage management database also addresses the security complication. Included in the database is information about which users can use or view which files. The storage management database prevents unauthorized users from

accessing files for which they have no security clearance. It tracks each user and ensures that only approved users have access to files stored on the backup media.



The Backup Management Program

Writing files to the backup media and managing the storage management database requires an application program, such as TapeWare. Two of the most important functions of TapeWare are *managing the storage management database* and *creating and running jobs*. TapeWare manages the storage management database to keep track of files and to ensure security. TapeWare also creates and runs jobs, such as backup and restore jobs which transfer files back and forth between backup devices (such as tape drives) and file servers and workstations.

Managing the Storage Management Database

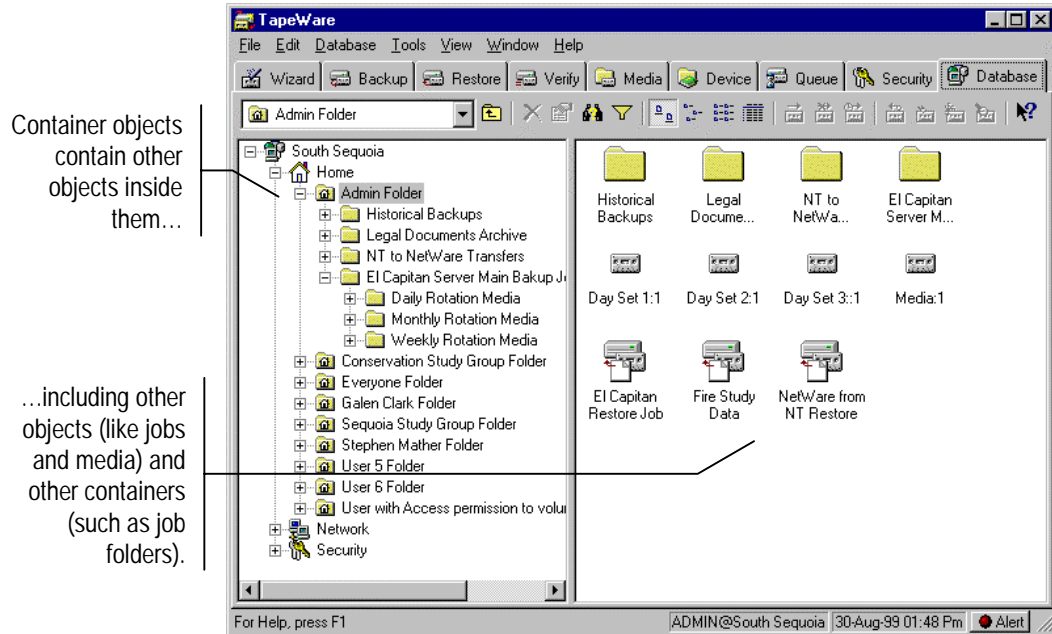
Much of the power and usefulness of TapeWare comes from its extensive capabilities to manage the storage management database. There are three important concepts associated with the database: *objects*, *properties*, and *storage management zones*.

Objects

The storage management database collects and organizes information about **objects**. An object is any file, machine, tape, or user about which TapeWare needs to store data. Examples of objects which TapeWare tracks in its database

include tapes, tape drives, network servers, occurrences of files, backup jobs, users, and so forth.

Objects that contain other objects within them are referred to as **containers**. A simple example of a container is a folder. It contains other objects within it, including objects that are not containers, such as jobs and files, and objects that are containers, such as other folders.



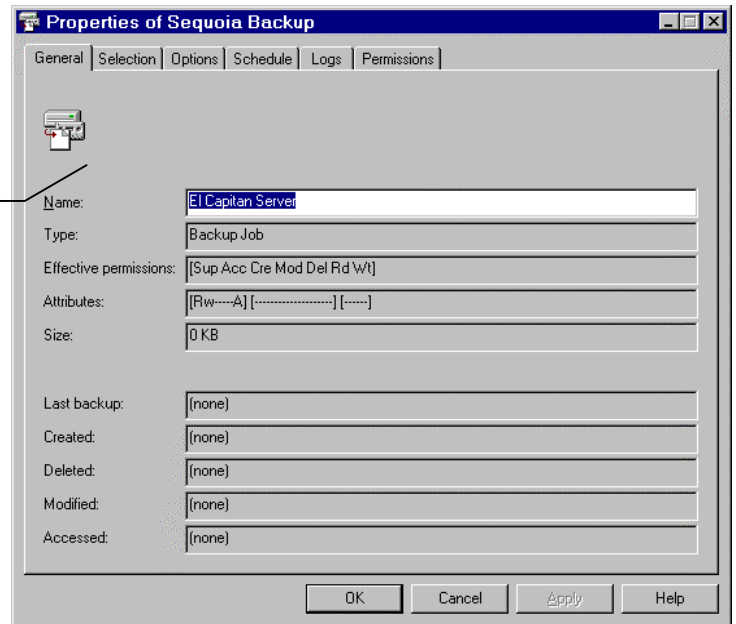
Properties

For further information on property sheets, see Chapter 12, "Objects and Properties Reference."

The information about each object that TapeWare stores in its database are called **properties**. The properties of each object include important information about that object, such as what kind of object it is, who has security clearance to use it, and its relationship to other objects.

For example, an individual backup tape is an object in the TapeWare database. Some of the properties of that tape stored in the storage management database include the name of the tape, when it was created, who has clearance to use the tape, and whether or not the tape can be erased.

You use property sheets, such as this property sheet of a backup job, to set, modify, and view the properties of an object.



Working with TapeWare objects and properties is easy. Even though when you work with TapeWare you work primarily with objects and properties, extensive knowledge about objects and properties is not necessary. If you know how to use the Windows® Explorer™, you already know most of what you need to use TapeWare. Part I of this manual will provide all the information you need to work with TapeWare effectively.

Storage Management Zones

TapeWare is not limited to just one database. On a large network, it may be useful to have multiple storage management databases, each addressing different storage needs. For example, there may be a separate database for each work group or department, even though they are all on the same network.

If your network has multiple TapeWare databases, then you choose which database you wish to use by selecting a **storage management zone** when you log on to TapeWare. Choosing a storage management zone is simply a way to select which database you wish to use.

Each TapeWare storage management zone is overseen by a **TapeWare administrator**. It is the TapeWare administrator's job to manage the security and integrity of the files in his or her storage management zone.

For further information, see Chapter 11, "Permissions and Security Reference."

Storage Management Zones and Security Storage management zones also help insure network security. TapeWare uses storage management zones to promote security in two ways: first, by preventing users from working with more than one storage management zone at a time and, second, by allowing file servers and workstations to be the member of *only* one storage management zone.

Note the following about machines, users, and storage management zones:

- Workstations or file servers (called a *machine container*) can only be the member of one—and only one—storage management zone. Because these machines belong to only one database, their peripherals, such as disk drives and backup devices, can belong to only one storage management database.
- Tapes or other media created in one storage management zone cannot be used in another storage management zone without following special procedures. This insures that there is no improper access to secure files and data.
- Each storage management zone must have at least one backup device, such as a tape drive. Further, that backup device can only belong to one storage management zone: it cannot be shared among multiple zones. (However, a storage management zone could have multiple backup devices.)
- From one workstation, a user can work in other storage management zones besides the zone his or her workstation belongs to. This means they can remotely administer jobs for other storage management zones besides the zone their workstation belongs to. They cannot, however, work in multiple storage management zones at the same time.
- A user can work from their home machine with other storage management zones besides the one to which their machine belongs. Their machine, along with its drives, peripherals, and accompanying data, however, always remain in a single storage management zone. This helps to prevent the unauthorized sharing of data between storage management zones. Thus while users can work outside of their database, the workstations and file servers they use always stay within their home zone.

Creating and Running Jobs

TapeWare creates backup tapes and restores files with **jobs**. Working with backup devices and the LAN, TapeWare jobs either back up network file servers and workstations onto media or restore files from media onto file servers or workstations. When you want TapeWare to backup or restore a file, you create and run a job.

There are several kinds of jobs, including backup, restore, and verify jobs. Every kind of job you create and run has six components: *creating the job*, *permissions*, *selection*, *options*, *scheduling*, and *running*.

Creating the Job You begin by creating a job, either a backup, verify, or restore job.

Permissions In order to create a job, you must have permissions to the objects that job will work with. For example, to create a backup job, you must have permissions to the tape drive, the tape, and the files you will backup. If you wish to backup files on another workstation, you must have permissions to that workstation and to the files on that workstation. Individual users are assigned permissions by the TapeWare administrator, who is responsible for ensuring the security and integrity of the backup system.

Selection Once you have permissions to a file, you must select it to be included in your job. You might select all files, only a few files, or perhaps only a single file. You select files first by selecting them in the selection window and second by having TapeWare ‘sort’ them with **filters**, which apply additional selection criteria such as date modified, type of file, and so forth.

Scheduling After selecting files, the job is scheduled to be run. A job may be scheduled to run later or to run immediately. It can be scheduled to run regularly or only once.

Options After scheduling the job, you specify the job’s option parameters. Some of the options you can specify include what backup device to use, what media to use, whether to automatically format the media, and so forth.

Running Finally the job is run. Many scheduled jobs are run automatically by TapeWare, but you can manually run a job at any time.

Every time you create a job in TapeWare, your job must include these six components. You begin by specifying a *type of job*. Before you can proceed further however, you must be certain that you have *permissions* to the objects, such as files and tape drives, that you wish to use. Your TapeWare administrator will work with you to determine which type of permissions you require. Then you *select* the files for your job, specify the *option* parameters, and finally *schedule* the job to run.

These six components of creating and running jobs are covered in detail in chapters 3 through 8 of this manual.

The TapeWare Workplace

TapeWare is designed to be easy to use. In fact, you may already know how to use many of its features. This chapter will familiarize you with all you need to know in order to use most of TapeWare's features, including keyboard shortcuts and mouse conventions.

In This Chapter

- | | |
|-----------------------------------|------------------------------------|
| • Installing TapeWare | • Property Sheets |
| • Starting TapeWare | • Menus |
| • The Logon Window | • Reports |
| • The Main TapeWare Object Window | • Working with the TapeWare Wizard |

Installing TapeWare

Installing TapeWare is simple and easy. The Install program automatically configures TapeWare to work on your machine with your network software and will prompt you when you need to specify installation parameters.

Before you install TapeWare, be certain you know the following:

- the name of storage management zone you wish to create.
- the machine on which you will locate the storage management database.
- which workstations or file servers will belong to the storage management zone.
- the Key Code for the product and configuration you have purchased.

If you are uncertain about creating a storage management zone and assigning machines to a storage management zone, see Chapter 1 and Chapter 10 of this

manual for more information. The following part discusses product configuration.

(If you will be working with NetWare's Btrieve database records, you must configure a special control file to backup and restore your database. See Appendix I for more information.)

Product Configuration

When fully implemented, TapeWare can work with any network configuration of Windows NT, Windows 95/98, NetWare, and DOS machines, with no limitation on the number of machines, servers, or backup devices.

However, what product you have purchased may limit your installation choices. The number of servers in a storage management zone and the network platform TapeWare runs on depend on which product you purchased. You can purchase license agreements for a single machine, a single server, or an unlimited number of servers. Similarly, your license agreement determines whether TapeWare will run only on a Windows NT network, on a NetWare network, or on both (including "mixed" networks).

If you have purchased a single server license agreement, you can have one NetWare or Windows NT file server in that zone. If you want to have multiple servers in a single storage management zone, you can purchase an upgrade to your license agreement.

If you have purchased a license agreement to install TapeWare on only one network platform, the Install program will only install TapeWare on the specified network platform. If you want to install TapeWare on another network platform, you can purchase an upgrade to your license agreement.

If you want to create more than one storage management database, you must purchase an additional software license. Each product purchased creates one and only one storage management zone.

Installation and Storage Management Zones

You should install TapeWare on the workstation or file server where the storage management database is going to be located *first*. When you do so, you create a storage management zone and specify the folder (the directory) in which the database will be stored. TapeWare needs this information when configuring other machines that belong to this storage management zone.

After first installing TapeWare on the machine where the storage management database is going to be located, later, when you install TapeWare on other workstations or file servers, TapeWare looks to see what storage management zone is available on the LAN. As part of the installation procedure, you will be

prompted to confirm that you want this machine to be part of the storage management zone you have already created.

Installation Instructions

(The following instructions assume that your CD-ROM drive has been assigned the drive letter D:. If this is incorrect, substitute the proper drive letter.)

DOS

Change the drive to the CD-ROM drive and then type **INSTALL**. Consider this example.

```
C:> D: [ENTER]
D:> INSTALL [ENTER]
```

NetWare

If you have a CD-ROM drive attached to the file server, mount the CD-ROM drive and type **LOAD TapeWare:INSTALL**. Consider this example.

```
SERVER: CD MOUNT cd-rom device [ENTER]
SERVER: LOAD TAPEWARE:INSTALL [ENTER]
```

If you do not have a CD-ROM drive attached to the file server, you must create an install directory and copy the install program to that directory. Consider this example.

From a workstation:

```
C:> MD F:\TWINS [ENTER]
C:> MD F:\TWINS\NET [ENTER]
C:> COPY D:\INSTALL.NLM F:\TWINS [ENTER]
C:> COPY D:\NET\*.* F:\TWINS\NET [ENTER]
```

Then from the server:

```
SERVER: LOAD SYS:TWINS\INSTALL [ENTER]
```

Windows (95/98 or NT)

If you are running Windows 95/98 or Windows NT 4.00, insert the CD in your CD-ROM drive. Setup.exe will automatically be run. (If this fails to work, double click the Setup.exe file on the CD).

If you are running Windows NT 3.51, change the drive to the CD-ROM drive and then type **SETUP**. Consider this example.

```
C :> D: [ENTER]
```

```
D :> SETUP [ENTER]
```

The TapeWare Installation Manager

After starting Setup.exe, the **TapeWare Installation Manager** window appears. From this window, you can perform a number of installation related procedures, including installing TapeWare for the first time, installing connector and email packages, and updating your license agreement.

Install TapeWare

Select this option to install TapeWare for the first time or to reinstall TapeWare. The Installation Manager will guide you through each required step of the installation process. In general, you should accept the Installation Manager's suggested folder (directory) locations.

Note that you can also use this option to reinstall a newer, updated version of TapeWare than your current version. If you are using Version 6.0 or later, the Installation Manager will keep your current storage management database.

Remove TapeWare

Select this option to uninstall TapeWare. (Note that you need not uninstall TapeWare before upgrading to a newer version.)

Install an Option

Select this option to install an optional package. These packages include a package to configure automatic email messaging as well as optional connector packages to work with such database programs as Microsoft Exchange and Microsoft SQL.

Some database connector packages are only available with the proper license agreement. You can purchase the appropriate upgrade by contacting Customer Service.

Remove an Option

Select this option to remove an optional package that you had previously installed.

Update License

You can upgrade TapeWare from your current version by purchasing a new license agreement. For example, you might purchase a new license agreement to backup additional file servers or to work with both NetWare and Windows platforms.

If you have purchased an upgrade, rather than reinstalling TapeWare in order to upgrade your software, you can more simply and quickly update your license agreement by using this option. Be certain to have your new key code available when you select this installation option.

Installing TapeWare as a Service

When TapeWare is installed and run as a service, it will automatically be run whenever the system starts up. When it runs as a service, it will run in the background *without the user interface*, although you can make TapeWare's user interface active at any time by double-clicking on the TapeWare icon. Later, when you close TapeWare, it returns to service mode and runs in the background again.

Installing TapeWare as a service can be very useful because TapeWare will run backup jobs automatically and unattended. This can be essential after a power loss, for example. If TapeWare is installed as a service and there is a power loss, after power returns, TapeWare is automatically run in the background when the system restarts, insuring that any scheduled jobs are run.

When you install TapeWare on a Windows 95/98 or Windows NT machine, the install program will ask you if you want to install it as a service. To install it as a service, check the appropriate check box in the Install program's **TapeWare Service** window.

Note TapeWare is only available as a service on Windows 95/98 and Windows NT platforms.

Installing Backup Devices

TapeWare will automatically recognize any backup devices attached to your file server or workstation, provided they are connected to the machine via a SCSI connection. If your machine recognizes the backup device as a valid SCSI device, then it will automatically appear as a backup device in the storage management database. This is true, as well, for autoloaders and devices with multiple tape drives.

If a device you expect to be available for use does not appear, check to see that your machine or network recognizes that device. After being certain that you have properly installed the device according to the manufacturer's instructions, open the Windows Explorer (or similar program) and check to see that the backup device appears as a device on the appropriate machine. If it does not appear, begin by checking the cable connection. Then run a utility, such as the Add New Hardware control panel, in order to make the device available to your machine. If this fails, contact the manufacturer of the backup device.

Note that TapeWare can only work with devices that are *both* recognized by your file server or workstation *and* attached via a SCSI or ATAPI/EIDE connection.

Starting TapeWare

After you have installed TapeWare, you start it like any other program.

To start TapeWare in Windows® 95/98 or Windows NT™, click on the **Start** button on the **Taskbar** and select **TapeWare** from the **Program** submenu. You can also create a shortcut to TapeWare and put it on the desktop. Alternatively, you can double click the Twadmin.exe file in the **Windows Explorer** window.

To start TapeWare on a NetWare® platform, run TWAdmin.NCF from SYS:SYSTEM. It is not necessary to specify a search path. Consider the following example:

SERVER: TWAdmin

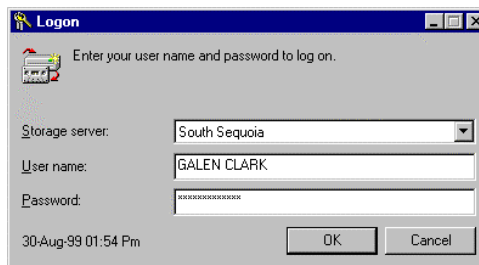
To start TapeWare on a machine running DOS, position to the directory where you installed TapeWare. Then type TWAdmin. Consider the following example:

C:\...>TWAdmin

The Logon Window

Each time you start TapeWare, you are shown the **TapeWare Logon** window.

The Logon window



To log on, you must select a storage management server, enter your user name, and enter your password.

Selecting a Storage Management Zone

When your TapeWare administrator set up TapeWare to run on your LAN, he assigned your workstation to a storage management zone. Your workstation, along with its drives and peripherals, can only be the member of one database

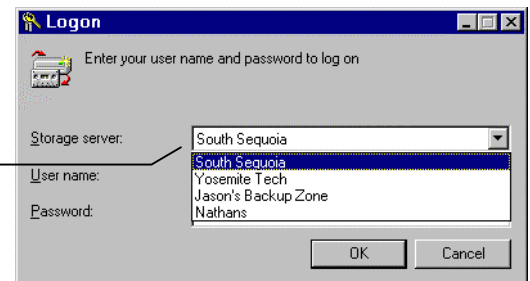
zone. The name of this zone is the *default* name that appears in the **Storage Server** list box.

Normally, you should leave the default name in the list box unchanged. This is because you will typically want to work with the storage management zone your workstation belongs to.

Occasionally, however, you may wish to work in a different zone. You might be asked by a coworker or your TapeWare administrator, for example, to run a TapeWare job in a different storage management zone.

To select a storage management zone other than the default zone, first click the arrow next to the **Storage server** list box. You will see a list of possible storage management zones. Select the storage management zone you wish to use.

To select a storage management zone, click on the **Storage server** list box and select the storage management server you wish to use.



Note Although you can log on to different storage management servers, you can only create and run jobs within a single storage management zone. Further, you can only access files and devices in a single storage management zone. This means, for example, that you will be unable to restore files backed up from workstations in one storage management zone to workstations in another storage management zone. (If you need to share data from one storage management zone to a different storage management zone, see “Import Media” in Chapter 9.)

User Name and Password

After selecting a storage management zone, type in your user name and password.

If you type your name or password wrong, you will be asked to reenter your name and password.

In order to log on, your TapeWare administrator must first assign you a user name and a password. If you are having difficulty logging on, ask your system administrator to specify again the exact spelling of the user name and password assigned to you.

Changing Your Password

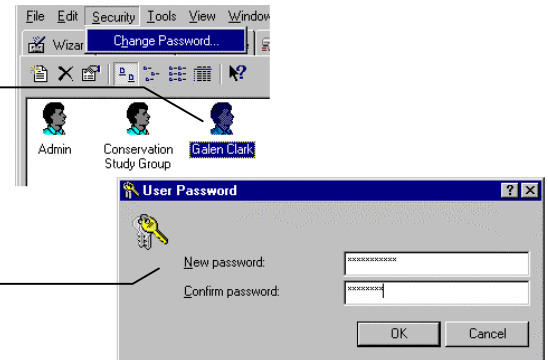
It is a good idea to regularly change your password, particularly if you are working with sensitive and important data. To change your password, make the **Security** tab active by clicking on it. Then choose **Change Password...** from the **Security** menu or the **Shortcut** menu. Change your password in the **User Password** window.

◆ To Change Your Password

1. Click on the **Security** menu to make it active.
2. Click on your **User Object** icon.
2. Open the **User Password** window by either
 - selecting **Change Password...** from the **Security** menu, or
 - clicking your user object with the right mouse button and selecting **Change Password...** from the **Shortcut** menu.
3. Type in your old password and then your new password.
4. Confirm your new password and then click **Ok**.

To change your password, highlight your **User Object** icon, and then select **Change Password...** from the **Security** menu

In the **New Password** window, type in your new password, then confirm password.



When selecting a password, remember that some passwords are notoriously easy to break. For example, because many persons use their birth date or name of their spouse, these are not good choices for passwords.

Tip The TapeWare administrator can change a user's password without knowing the user's current password. When the TapeWare administrator selects **Change Password...** from the **Security** or **Shortcut** menu, TapeWare does not require that the old password be entered before changing the password. This is useful when the user has forgotten his or her password.

Logging on the First Time

The default user name for the TapeWare administrator is **ADMIN**. There is no default password required for this user to log on.

Warning TapeWare administrators have unlimited access to all of the objects in the database. Any user who logs on as the TapeWare administrator will have complete access to all of the files and machines on the database.

Your first security step should be to change the TapeWare administrator's password. Change your password on the **Security** tab. Select your User object and select **Change Password...** from either the **Security** menu or the **Shortcut** menu. Before continuing with the rest of this section, be certain that you have changed your password

Grace Logons

Your TapeWare administrator may have set up your password to expire after a set period of time. For example, your password may be invalid after 60 days. This forces you to change your password regularly.

If your password has expired, TapeWare will prompt you to change your password. If you choose not to change your password, TapeWare may still let you log on, even with an expired password. Logging on with an expired password is called a **grace logon**. Your TapeWare administrator will determine how many grace logons you are allowed.


When your password has expired and you have used all of your grace logons, TapeWare will ask you to change your password when you log on.

Logging Out

On occasions you may want to change the database zone you are using or want to log on as a different user. Although you could quit TapeWare and restart the program, it is quicker and faster to log on again without quitting TapeWare.

To log on again, choose **Log Out** from the **File** menu. You will be presented with the **TapeWare Logon** window and asked to log on again.

Select **Log Out** from the File menu to log on to a new zone or to close the main window.



Logging Out and Running Scheduled Jobs

For further information, see "Running Scheduled Jobs," Chapter 8.

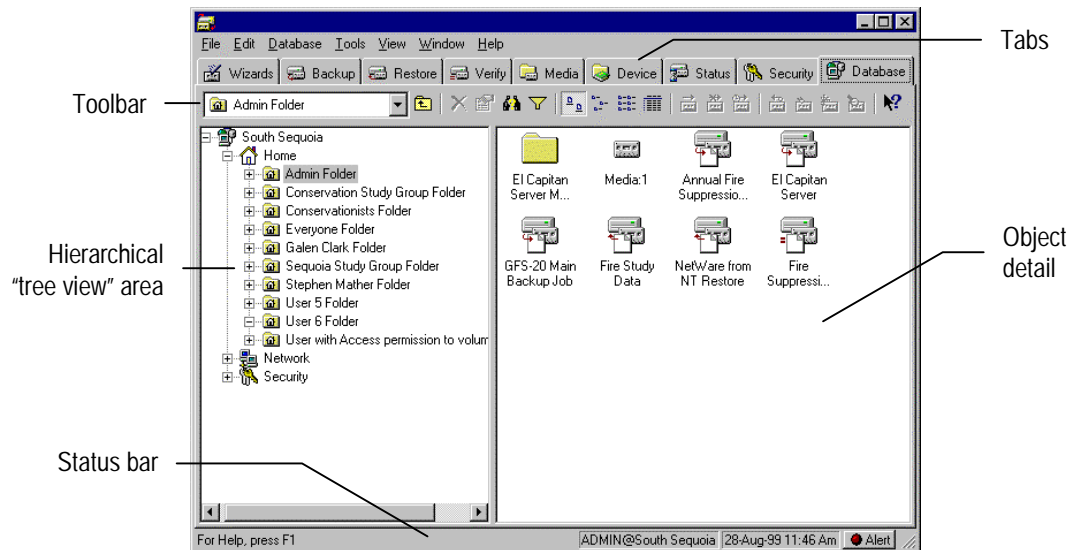
TapeWare can run scheduled jobs even when no one is logged on. (Only *scheduled* jobs can be performed when no one is logged on to TapeWare.)

When you leave your workstation, you may need to leave TapeWare open. In order to not let any other users have unauthorized access to the LAN, log out of TapeWare before leaving your workstation. Any scheduled jobs will still run, but no unauthorized users will be able to work with TapeWare unless they can log on.

The Main TapeWare Object Window

You use the main TapeWare object window to view, create, and manipulate TapeWare objects, such as jobs and tapes.

In addition to the menu bar, the TapeWare window has the following parts: tabs, toolbar, the hierarchical or "tree" view area, the object detail area, and the status bar.




Tabs

At the top of the window are eight **tabs** which help to group and organize similar sets of objects. For example, you use the **Backup** tab to display and work with backup jobs, the **Restore** tab to display and work with restore jobs, and so forth. To display a different tab, click on the tab you wish to view. You can also display a different tab by using the **View** menu.

Toolbar

The **Toolbar** has a list box and several buttons. Not all buttons are on all tabs and some buttons on a tab may be unavailable. When a button is grayed out, it is not available because no object that it can work with is currently selected.



- The list box contains the name of the folder (or container) currently displayed in the object detail area.
- The **Up One Level** button changes the currently displayed folder (or container) to the next higher container in the hierarchy, that is, the container that contains the current container.

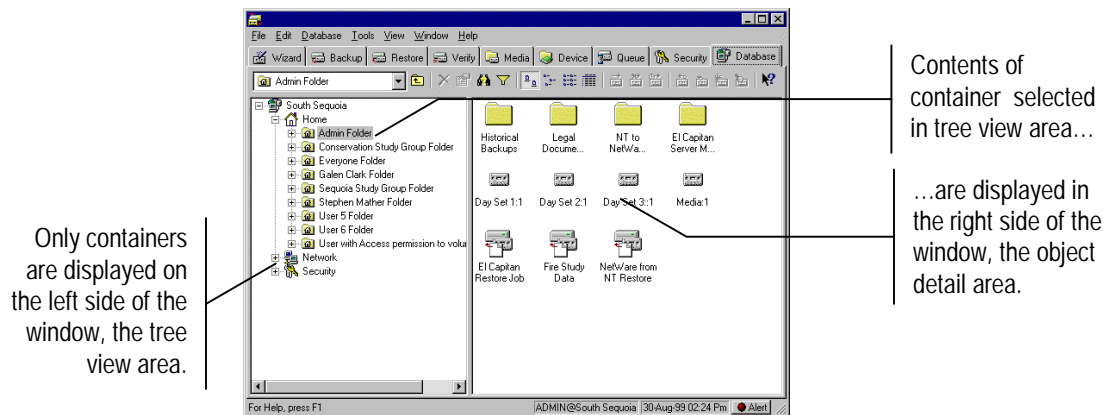
- The **Large Icon**, **Small Icon**, **List**, and **Details** display buttons determine how objects are displayed in the object detail area.
- The **Properties** button displays the properties of any object you have selected in the object view area.

You can use the **View** menu to either hide or show the **Toolbar**.

The Object View Areas

The left side of the window displays containers, such as folders. Examples of other container objects include networks, workstations, and drives. These are displayed in hierarchical or “tree” view. This portion of the window is referred to as the **tree view area**.

The right side of the window is referred to as the **object detail area**. It displays the contents of the container currently selected in the tree view area. If you select an object in the tree view area, you will see its contents displayed in the object detail area on the right side of the window.



The object view area is designed to be easy to use and if you have worked with Windows Explorer, you probably already know all you need to work with it. If you are new to working with objects presented like this, keep the following in mind:

- The left side of the window, the tree view area, only displays *containers*, that is, objects that hold other objects. Objects that do not contain other objects are displayed only on the right side of the window, in the object detail area.
- To display the objects in a folder or other container, open it by clicking on it in the tree view area.



Expand tree icon.

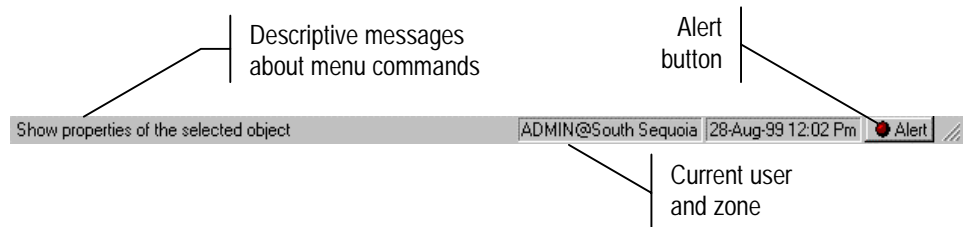
- To view the folders within a folder, expand the “tree” by clicking on the expand tree icon next to the folder. Alternatively, you can double click the folder, and it will both expand and display its contents in the object detail area.

Note You can instruct TapeWare to only show the expand tree icons when a container contains other containers within it. Select **Preferences...** from the **File** menu and check the **Smart expandable indicators** check box. Before displaying the expand tree icon, TapeWare will check to see if that container contains other containers in it. If it doesn't, the expand tree icon will not be displayed.

There are numerous keyboard shortcuts available that make it easier to work with objects in TapeWare. For more information, see “Keyboard Shortcuts” later in this chapter.

The Status Bar

The status bar displays short descriptive messages about the menu commands on its left side. The middle of the status bar displays the current user's name and the database zone they are currently logged on to.



On the right side is the **Alert** button. This button flashes whenever there is a problem that requires your attention. For example, TapeWare might not be able to locate a backup device you specified for a job and so it will send an alert to the **Alert** window notifying you of the problem. When you click on the **Alert** button, TapeWare brings up the **Alert** window so you can view any pending alerts.

You can use the **View** menu to either hide or show the **Status Bar**.

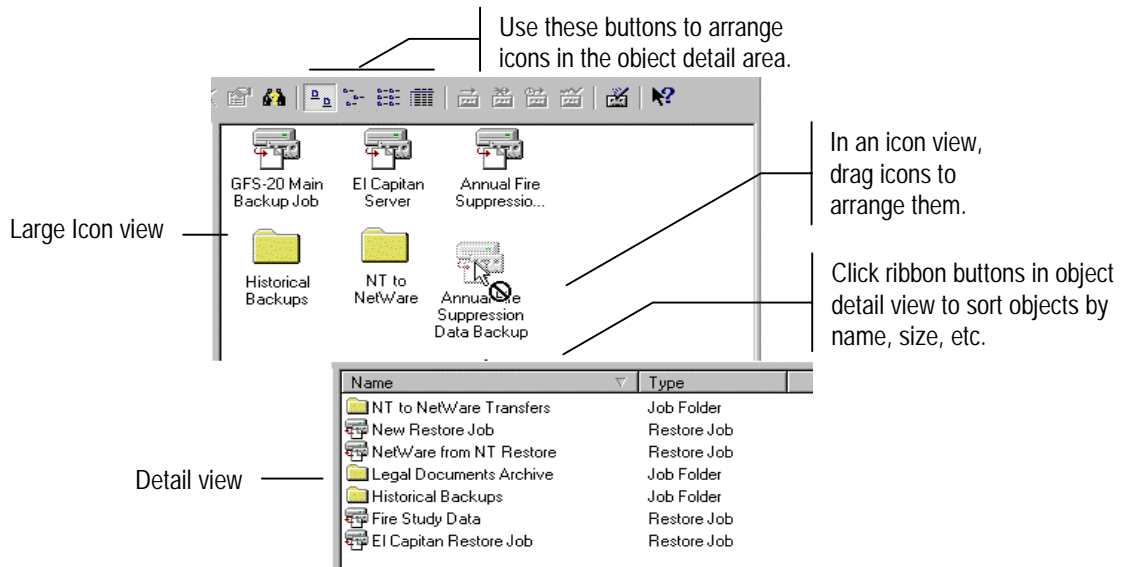
Working with Objects in the Main TapeWare Window

You can easily change how objects are displayed in the main TapeWare object window. This will assist you in working more effectively and quickly.

Arranging Objects

In the object detail area, you can arrange the objects in several different ways using either buttons on the **Toolbar** or the **Arrange Icons** submenu on the **View** menu. This will allow you to sort or arrange the icons either according to name, date, or type of object. You can also

- drag the icons to arrange them when working with either large or small icons; and,
- use the ribbon at the top of the object detail area to change how the objects are displayed when working in detail view.



Keyboard Shortcuts

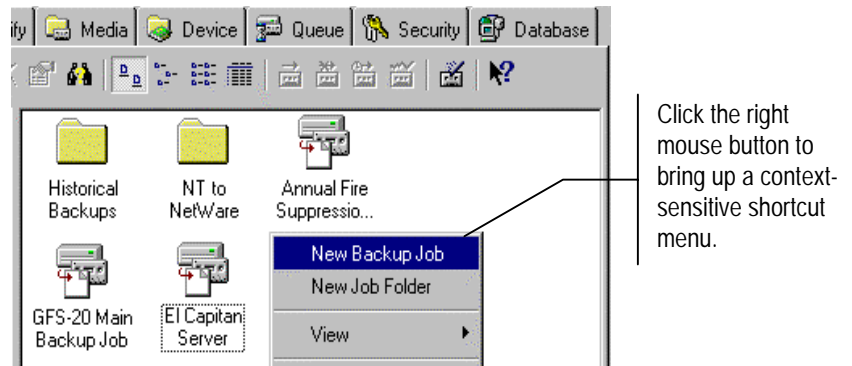
Although you may find using the mouse to be the most intuitive way to work with objects in the main TapeWare window, there are several keyboard shortcuts that will help speed up your work. The next time you work with TapeWare, try these keyboard shortcuts.

- The TAB key will move the active or highlighted area to a different area of the window.
- The PLUS SIGN (+) on the numeric keypad or the RIGHT ARROW key *expands* the tree in the tree view area;
- The MINUS SIGN (-) on the numeric keypad or the LEFT ARROW key *collapses* the tree in the tree view area;

- The ASTERISK key (*) on the numeric keypad will *expand* the tree as far as possible.
- The ARROW keys also select objects in the tree view area and in the object detail area, as well as change the active tab.

Shortcut Menus

In most windows, when you click the right mouse button, the **Shortcut** menu appears. The shortcut menu list commands that pertain to the particular portion of the screen you clicked on. This is often the fastest and easiest way to create new objects and modify existing objects.



The Find Button



The Find button

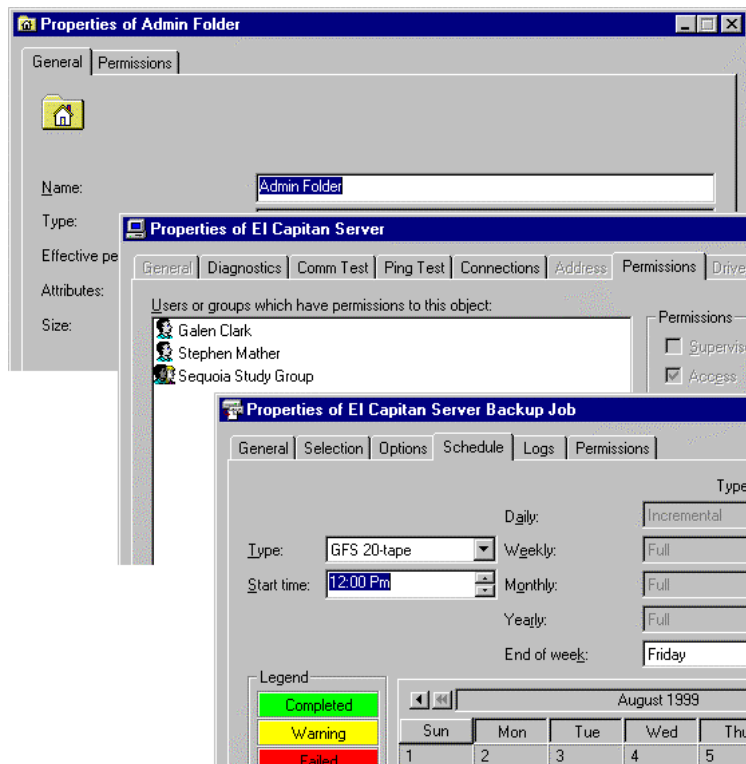
The **Find** button lets you locate objects in the database quickly without searching through multiple volumes and directories.

To use the **Find** command, click on the **Find** button, select **Find...** from the **Edit** menu, or press CTRL+F. In the **Find** window, type in the name of the object you wish to find. TapeWare will search through the database, attempting to locate that object. When the object is found, it will be displayed and highlighted.

Note that the **Find** command is not case sensitive; additionally, you may use the wildcard characters ? and *.

Property Sheets

Every object in the TapeWare database has a **property sheet** associated with it. Each property sheet has two or more tabbed pages on it that display the properties of that object.



Examples of property sheets for a job folder,...

...a network server,...

...and a backup job.

Opening Property Sheets




The Properties button.

You can display the property sheet of an object in one of three ways.

- Select the object with the mouse or keyboard, and then click the **Properties** button on the **Toolbar**.
- Click the right mouse button on the object to display a shortcut menu, then select **Properties**.
- Select the object and press ALT+ENTER.

You can view the properties of an object either by selecting it and clicking the Properties button on the toolbar...



...or selecting Properties from the shortcut menu.

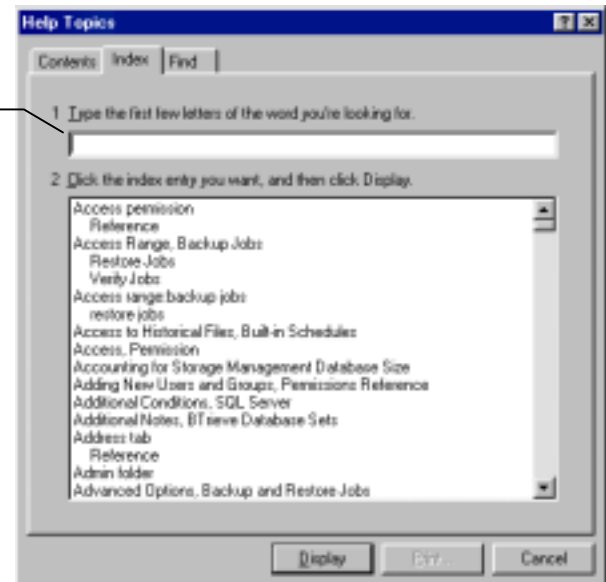
Note Property sheets work just like other windows, except that their size cannot be changed. You can leave them open when you return to working in the main TapeWare object window and you can have multiple property sheets open at once.

Menus

You will find the menus easy and intuitive to use. Most of the menu commands are discussed in some detail in later chapters of this book.

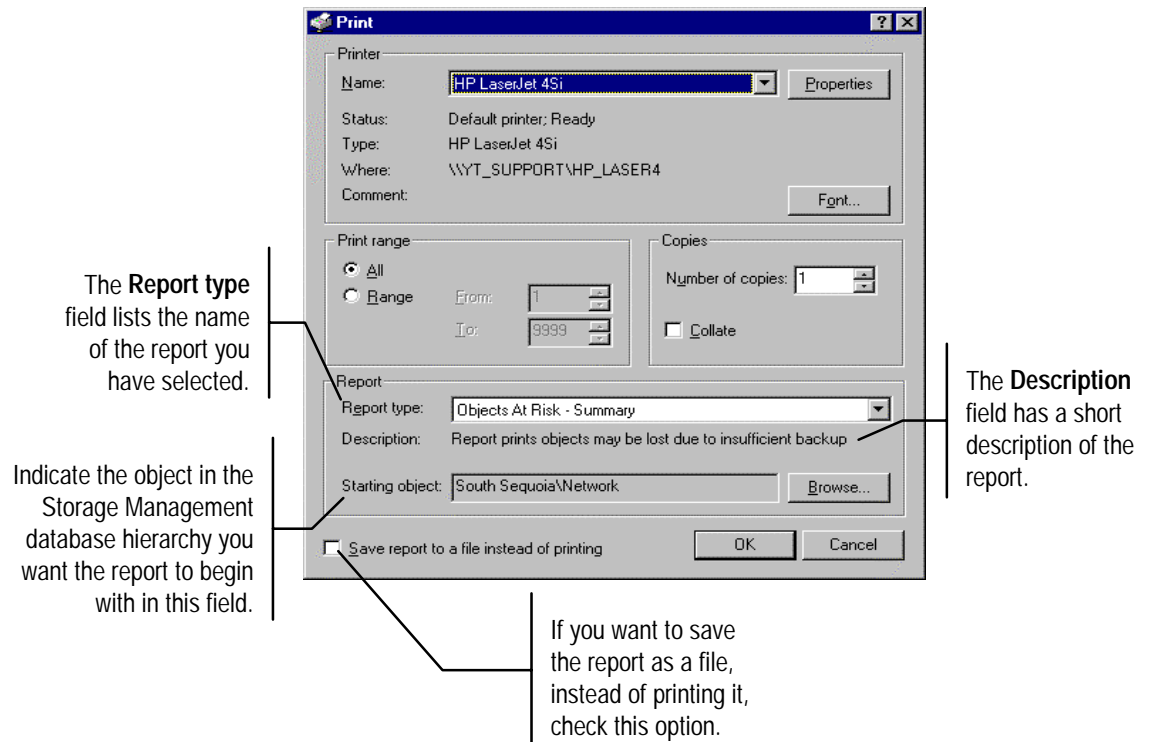
If you are having difficulty with a particular command, look in the index of this manual under the name of the menu. The **Status Bar** also displays a short description of the highlighted menu command. Alternatively, select **Help Topics** from the **Help** menu. Then type the name of the command with which you are having difficulty in the list box on the **Index** tab of the **TapeWare Help** window.

You can get help on any menu command by typing the name of the command on **Index** tab of the **Help Topics** window.



Reports

TapeWare has a series of helpful diagnostic and summary reports available. To run a report, select **Print report** from the **File** menu. In the **Report type** list box, you will see a list of the reports available. When you select a report, a short description appears in the **Description** field.



Before printing a report, specify which objects you want the report to cover by selecting a **Starting object**. The report will generate information for all of the objects in the TapeWare hierarchy that are below the object specified in the **Starting object** field. You can change the starting object by clicking on the **Browse** button and selecting a new object.

You can also specify the font the report will be printed in by clicking on the **Font...** button.

These reports can be very useful. For example, the "Objects at Risk" reports will tell you which files, directories, and volumes have not been sufficiently backed up and are therefore at risk unless additional backups are performed. The "Database Listing" reports gives you a detailed report on particular objects in the storage management database. By changing the starting object, you can customize this report to provide you with data only about particular objects, such as users or media. For example, for detailed information about Users and Groups, set the **Starting object** setting to the **Security** folder.

Additional Report Capabilities

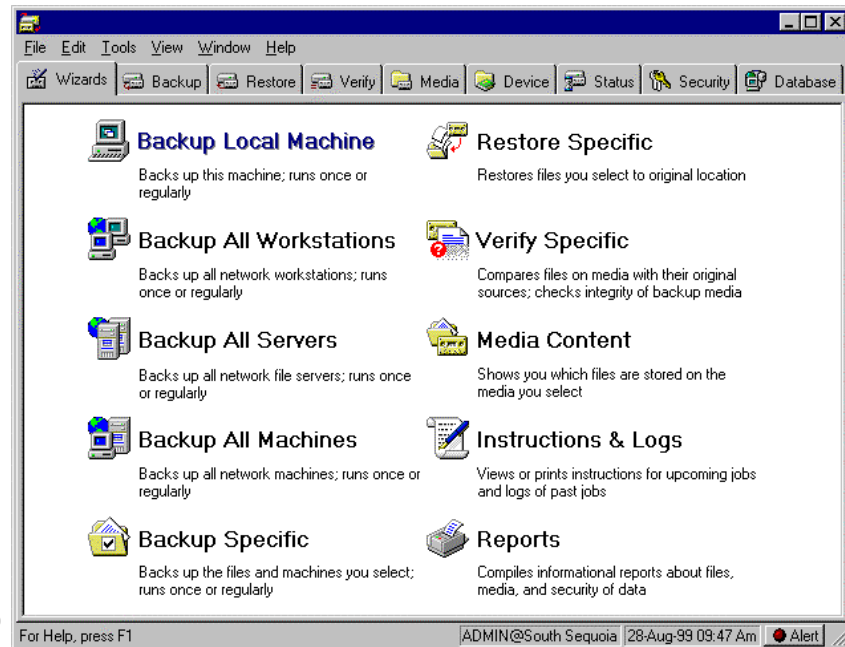
Besides its reporting capabilities, TapeWare has several advanced features to help you track and compile necessary information.

- The **Instructions** command on the **Tools** menu lets you check to see which media and devices are required for the scheduled jobs. See “Viewing and Printing Job Instructions” in Chapter 8 for more details.
- The **Email** attachment allows you to email the logs of jobs to any valid email address. For more information, see “Email Support for Job Logs,” in Chapter 8.
- You can also create *audit logs* for files, directories, and volumes, in order to monitor the backup and restore activity for these objects. See “Viewing and Printing Audit Logs” in Chapter 8 for additional information.

Working with the TapeWare Wizard

The TapeWare Wizard gives you a fast and quick way to perform most tasks. You can use the TapeWare Wizard to create and schedule backup jobs, restore jobs, and verify jobs. You can also create and edit reports from the **Wizard** tab. The TapeWare Wizard will guide you through each step necessary to create and run a job or create a report.

The **Wizard**
tab



The name of each TapeWare Wizard option indicates what task that wizard helps you perform. The descriptions are self-explanatory. The backup, restore, and verify options create jobs; for additional information about creating these jobs without the TapeWare Wizard or about the jobs the TapeWare Wizard created, consult Chapters 3 through 8 of this manual. The **Media Content** wizard is a specialized implementation of the **Query Filters** window, discussed in greater detail at the end of Chapter 9 and in Chapter 12. The **Instructions and Logs** wizard tells you which media and devices are required for scheduled jobs and results of previously executed jobs; for additional information, consult “Viewing and Printing Scheduled Job Instructions,” Chapter 8.

When an option is grayed, that option is not available due to your license agreement. For example, some license agreements do not let you backup file servers. If you want to backup a file server or more than one file server, you can purchase an upgrade to your license agreement. This upgrade will enable these grayed-out options.

Creating TapeWare Jobs

To transfer files back and forth from backup devices, such as tape drives, to network workstations and file servers, you create and run **jobs**. You organize and store these jobs in **folders** you create on **job tabs** in the main TapeWare object window.

In This Chapter

- Backup, Restore, and Verify Jobs
- Creating New Jobs
- Renaming, Deleting, and Moving Jobs
- Organizing Jobs with Folders

Overview

You use backup jobs in order to protect against loss of data due to disasters or equipment malfunction, to archive important files, and to create permanent historical records. Restore jobs allow you to transfer stored files on media back to file servers and workstations. Verify jobs compare the version of a file stored on media, such as tapes, with current versions of the file stored on machines on the LAN.

You can create folders to store these jobs in on one of the job tabs. These folders and the jobs stored within them can be renamed, deleted, and moved to new locations.

Backup, Restore, and Verify Jobs

TapeWare can create and run three types of jobs: backup jobs, restore jobs, and verify jobs.

Backup Jobs

Backup jobs copy selected files *from* workstations and file servers *to* various media, such as tapes. This media can then be stored, preserving a copy of the file for future use. You might create and run a backup job for one of the following reasons:

- *To insure the integrity of data should a disk drive on a workstation or file server fail.*

This is perhaps the most common type of backup job. Its purpose is to protect valuable information in case of a disastrous data loss. This type of backup allows a company or organization to return to work quickly, even after the failure of a main file server. To be effective, these backup jobs must be run regularly (normally daily) *without exception* in order that recent changes to files can be safely restored. This type of backup job is for *disaster protection*.

- *To remove files from a workstation or file server that are infrequently used.*

Some files are important to keep, but are never or rarely used. For example, you may wish to keep a copy of correspondence from last year for legal reasons, but have no regular need to access these files under normal circumstances. By backing up the files onto a tape or other media, you safely store the media, preserving a copy of the file, and then delete the file from the workstation or file server. TapeWare will keep track of which files you have backed up and which tape they are located on. As long as the media is undamaged and safely stored, you will be able to retrieve the file if necessary. This type of backup job is called an *archive job*.

- *To store a copy of a particular historical version of a file.*

Sometimes you may wish to keep a permanent record of a particular version of a file. For example, you may need to preserve a copy of company records as they exist on a certain date or before they are converted for use in a new program. You can store a copy of the file as it exists on a certain date and instruct TapeWare to make certain that this file and media it is on are not overwritten with other data. TapeWare will keep track of the file and the media in its database and you will be able to retrieve it if necessary. Unlike an archive job, the file that was backed up is not deleted from the file server or workstation. This type of backup job is sometimes called a *historical backup*.

Restore Jobs

Restore jobs copy files *from* backup devices *to* workstations and file servers. You might create and run a restore job when files on a workstation or file server

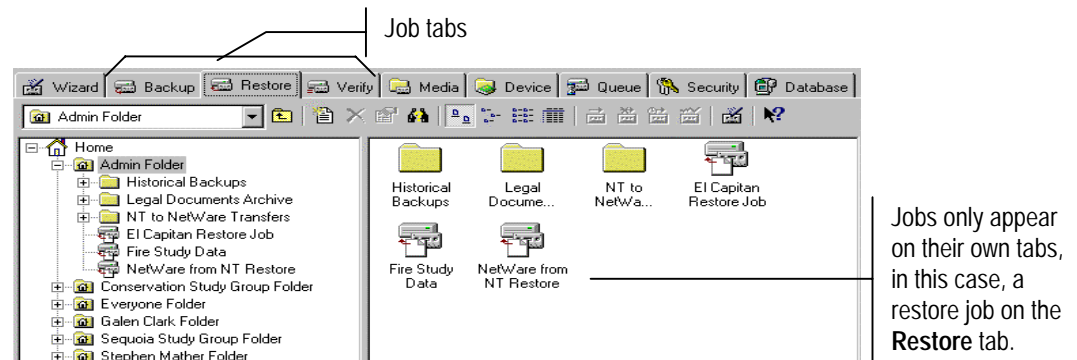
have been lost because of a disk crash, when you need view a file that has been archived (backed up onto a tape and then deleted), or when you need data from a particular historical version of a file.

Verify Jobs

Verify jobs compare a file on some media, such as a tape, with a file at a workstation or file server. These jobs *verify* that the two files are in fact the same file. A verify job is useful when you wish to be certain that a particular file, such as a program file, has not been corrupted or modified.

The Job and Database Tabs

The three different job types each have their own tab. You create, modify, and run backup jobs with the **Backup** tab selected, restore jobs with the **Restore** tab selected, and verify jobs with the **Verify** tab selected.



Note that a job only appears on its own type of job tab. For example, backup jobs are displayed on the **Backup** tab, but not on the **Restore** or **Verify** tabs.

However, you can also view all three types of jobs on the **Database** tab. But because the **Database** tab keeps track of all of the objects in the TapeWare database, it may appear very cluttered. Normally, when working with jobs, make the appropriate job tab active.

Creating New Jobs

There are three ways to create new jobs: with the wizard, from the job tab, and by copying old jobs.



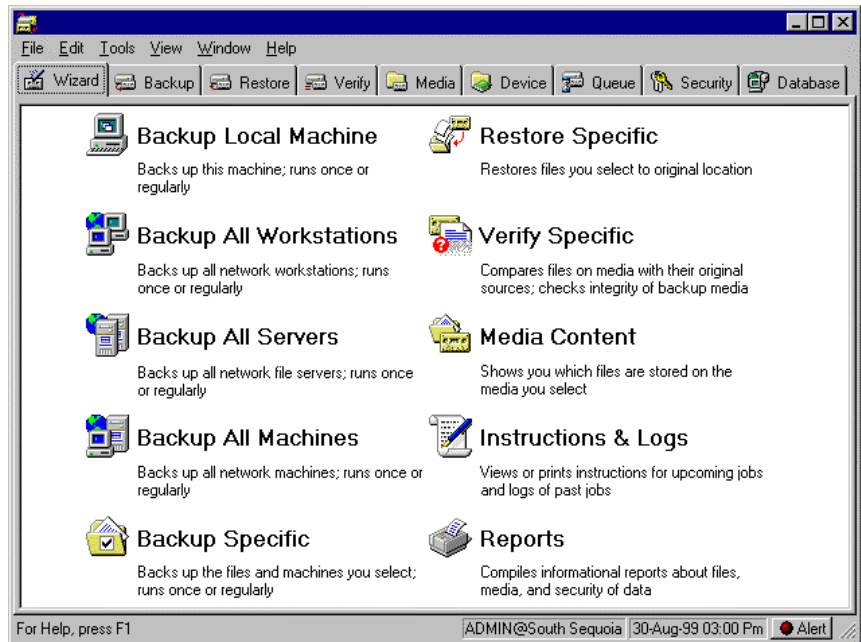
Wizard button.

Creating Jobs with the Wizard

The TapeWare wizard will help guide you through all of the steps necessary to create and run either a backup, restore, or verify job. This is often the fastest way to create a new job, especially when you are inexperienced using TapeWare. After you answer a few questions, the TapeWare wizard will create the job for you. You can then work with this job in the database just like any other job.

The fastest way to activate the TapeWare wizard is to click on the **Wizard** tab. Then click on the appropriate button to create either a new backup, restore, or verify job. You can also create reports from the **Wizard** tab which will help you monitor how jobs have run and the backup status of various files and databases.

Use the **Wizard** tab to create new backup, restore, and verify jobs.



You can also create a new job with the TapeWare wizard by either

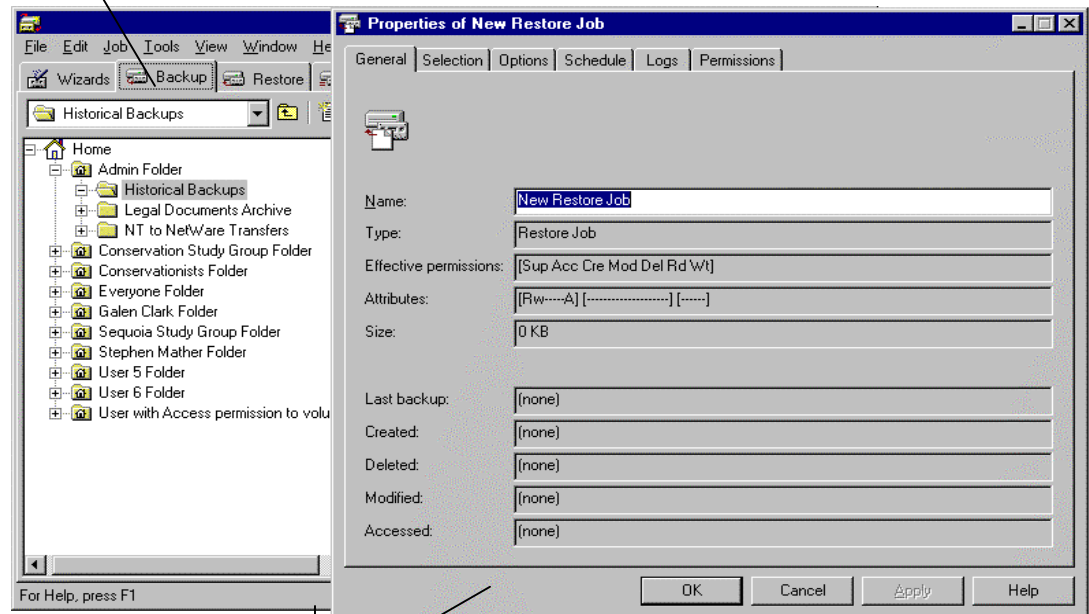
- selecting **Wizard...** from the **Job** menu,
- typing **CNTL + W**, or
- clicking the **Wizard** button on the **Toolbar**.

Creating New Jobs from a Job Tab

Commonly, you will create jobs while working with one of the three job tabs in the main TapeWare object window. When you create a job this way, TapeWare opens the property sheet of the new job so you can name the job, select files, and

schedule it to run. Selecting files and scheduling a job to run is covered in detail in chapters 5 through 7.

The type of job created depends on which job tab is active.



When you create a new job, its property sheet automatically opens



New Object button.

Which type of job you create depends on which job tab is active. Where TapeWare keeps track of the job depends on which folder is open in the tree display area. For example, if the **Backup** tab is active, TapeWare will create a new backup job, and so forth for the **Restore** and **Verify** tabs. If a personal or workgroup folder is open, TapeWare will store the job in that folder; if the **Everyone** folder is open, TapeWare stores the new job in that folder.

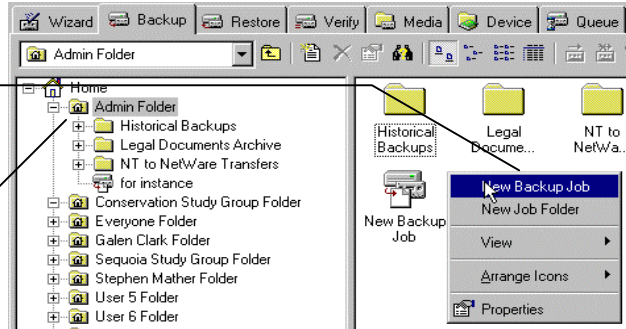
◆ To Create a New Job from a Job Tab

1. Make either the **Backup**, **Restore**, or **Verify** tab active by clicking on it. This determines what type of job TapeWare will create.
2. Open the folder in the tree view area where you want to store the job. The contents of that folder are displayed in the object detail area. TapeWare will store your new job here.

2. Create the new job by either
 - selecting **New...Job** from the **File** menu, or
 - clicking the right mouse button in the TapeWare object detail area and selecting **New...Job** from the shortcut menu, or
 - click the **New Object** button on the Toolbar, and select the appropriate job from the **New Object** window.
3. Type in the name of the new job in the **Name** box.

To create a new job, right click the mouse and select **New...Job** from the **Shortcut** menu.

New jobs are stored in the folder open in the tree view area.



Creating New Jobs by Copying

Creating a new job by copying an existing job is sometimes a useful method of creating a job. In particular, copying an existing job is appropriate when you want your new job to be like the old job in every way except for a few minor changes.

◆ To Create a New Job by Copying an Existing Job

1. Copy the existing job you wish to duplicate by either
 - selecting the existing job (with the mouse or keyboard) and pressing CTRL+C, or
 - clicking the right mouse button on the existing job and selecting **Copy** from the shortcut menu, or
 - holding down the CONTROL key as you drag an existing job to a new location, or
 - selecting **Copy** from the **Edit** menu.

2. Open the folder you want the new job to be stored in by selecting it in the tree view area. (To store the job in the same folder as the existing job, skip this step.)
3. Click the right mouse button where you want the new job to be pasted, and select **Paste** from the **Shortcut** menu. Alternatively, highlight the location where you want the job pasted and select **Paste** from the **Edit** menu.
4. Change the name of the new job.

Renaming, Deleting, and Moving Jobs



Delete button.


You can change the names of jobs or move them to new folders. Or, if you don't plan to use a job again in the future, you can delete it.

You rename, delete, and move jobs by either

- selecting the appropriate command from the **Edit** menu, or
- clicking the right mouse button on the job and selecting the appropriate command from the shortcut menu.

To move a job, drag it to a new location. You can also use the DELETE key and Delete button to delete jobs.

Right click on an object to bring up the Shortcut menu to Copy, Move, Delete, and Rename objects.



Note When you run a backup job, TapeWare uses its storage management database to keep track of the files you have backed up and the name of the tape on which they are stored. Deleting a job does not effect how the storage management database tracks files and tapes. TapeWare continues to track these files and tapes even after the job that created them has been deleted.

Organizing Jobs with Folders

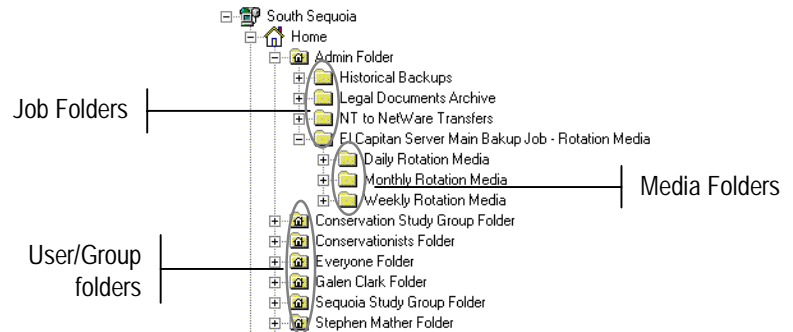
Every backup, restore, and verify job is stored by TapeWare inside a folder. You can create new folders to help organize your jobs or you can use existing folders.

When your TapeWare administrator added you as a user to TapeWare, he or she also created a personal folder for your use. Generally, because there may be many users in your TapeWare storage management zone, it is a good idea to store your personal jobs in your personal folder. Alternatively, store workgroup jobs in the appropriate group folder.

Types of Folders

There are three types of folders: **User/Group folders**, **Job folders**, and **Media folders**. Each of these folders is a container, that is, they store other objects within them. They differ from each other according to the type of object that can be stored within them.

- *User/Group folders* can only be stored in a special folder, called the **Home** folder. These folders can have either Job folders or Media folders within them; additionally, you can store jobs or media “loose” in these folders.
- *Job folders* can only be stored in User/Group folders or in other Job folders. These folders usually have jobs stored within them, although you can also store additional job folders within them.
- *Media folders* can only be stored in User/Group folders or in other Media folders. These folders usually have media stored within them, although you can also store additional Media folders within them. Media folders are discussed further in Chapter 9, “The Media, Device, and Database Tabs.”

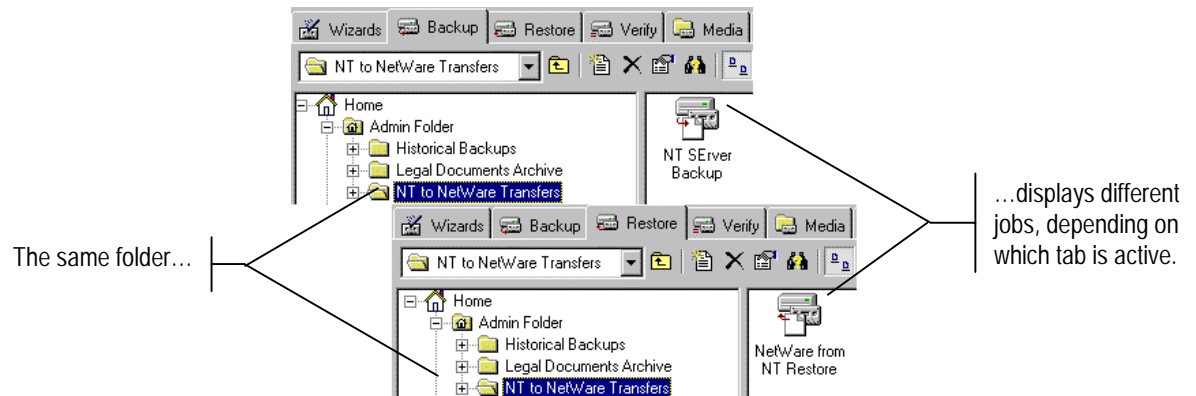


Folders and Job Tabs

Job folders are unlike jobs in one important way: Job folders appear on all the job tabs and on the database tab. Jobs, on the other hand, only appear on one tab (backup jobs on the **Backup** tab, and so forth) and on the **Database** tab. When you create a new job folder, that folder appears on all the job tabs and on the database tab.


Similarly, User/Group folders appear on every job tab, as well as the **Media** tab and the **Database** tab.

Media folders, however, only appear on the **Media** and **Database** tabs.




Because every folder is seen on every job tab, you may not need or want to have separate folders for different types of jobs. For example, a folder that stores only backup jobs will appear on all three job tabs, including the **Restore** tab and **Verify** tab. For this reason, it is better to name the folder after your workgroup or yourself, than by the type of job stored in it. Again, a single folder can store all of your personal or workgroup jobs.

Don't name job folders after type of job, since they appear on every tab.



The Home, Admin, and Everyone Folders

Three folders play a special role in every database: the **Home** folder, the **Admin** folder, and the **Everyone** folder.



The **Home** folder is the folder that stores all of the User/Group folders within it. It is always at the top of the hierarchy in the tree view area of a job tab. You are not allowed to store jobs “loose” in the **Home** folder, only inside other folders.

The **Admin** folder is a special folder used by the TapeWare administrator. Normally, only the TapeWare administrator has permissions to the **Admin**

folder. If you don't see it inside your **Home** folder, this is because the TapeWare administrator has not given you permissions to view it.

The **Everyone** folder is a folder to which every TapeWare user has permission. Your TapeWare administrator may place jobs in this folder to which he wants everyone to have access.

The Everyone Folder and Permissions

For further information on permissions, see Chapter 4.

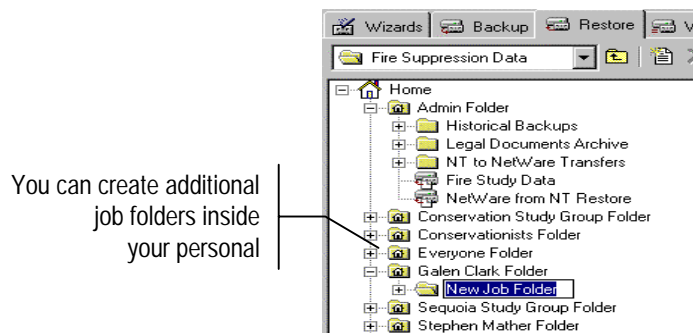
Because of the way TapeWare assigns permissions to new objects, if you create a new object inside the **Everyone** folder, normally everyone will have at least some permissions to it. For example, if you create a new backup job and store it in the **Everyone** folder, it is likely that every TapeWare user will have permissions to that folder and thus to the job.

To restrict the permissions of other TapeWare users to a job or folder, the best strategy is to use a folder to which only you or your workgroup have permission. Granting permissions are covered in detail in the Chapter 4, "Permissions."

Creating Job Folders

Usually, the best strategy for organizing your personal jobs or the jobs of your workgroup is to create a special folder in which to store them. This is particularly true because it makes it easier and quicker to manage permissions to those jobs.

When your TapeWare administrator adds you as a user to the database, he or she will create a folder for your personal use or the exclusive use of your workgroup. However, if you have the proper permissions, you can also create as many additional new job folders as you need and organize them in whatever way is convenient. You can create additional job folders either inside your personal or workgroup folder or inside the **Everyone** folder.



◆ To Create a New Folder

1. Select the existing folder in which you want to store the new job folder. (It cannot be the **Home** folder.)

2. Create the new folder by either
 - selecting **New Object...** from the **File** menu and then **New Job Folder** from the **New Object** window, or
 - clicking the right mouse button in the TapeWare object detail area and selecting **New Job Folder** from the shortcut menu.
3. Type in the name of the new folder on its property sheet.



New Object button.

You can also create a new folder in one step by clicking on it with the right mouse button in the tree view area and selecting **New Job Folder** from the shortcut menu. Alternatively, click the **New Object** button and select **New Job Folder**. TapeWare will create the new folder inside the folder you clicked on.

Tip After creating a new folder, be certain to specify which users have permissions to it. This is the simplest and fastest way to assign permissions to multiple objects stored in the folder.

Moving, Renaming, and Deleting Folders

You can move, rename, or delete a folder just like you can any other TapeWare objects, such as jobs; however, you cannot copy folders. When moving, deleting, or renaming folders, keep the following in mind:

Moving Folders Contents of a folder move with the folder to the new location. This may change the permissions to the objects stored in that folder.

Renaming Folders Only the name of the folder is changed. TapeWare still treats that folder and any objects associated with it in the same manner.

Deleting Folders Deleting folders also deletes their contents, including any other folders or jobs contained in that folder. Before deleting a folder, be certain that you wish to delete all of its contents.

Warning Once a folder has been deleted, its contents cannot be recovered. Be certain either that the folder is empty or that you no longer need the contents of the folder before deleting it.

Modifying Folders

You can rename, delete, and move folders using one of these methods:

- selecting the appropriate command from the **Edit** menu, or
- clicking the right mouse button on the job and selecting the appropriate command from the shortcut menu, or

- clicking and dragging the folder to a new location.

Permissions

Security is an important issue when managing a LAN. One of the most important functions of the TapeWare storage management database is to handle security. The storage management database prevents unauthorized users from working with objects to which they have not been granted security clearance. To ensure only authorized users can access sensitive data, TapeWare tracks the **permissions** of each user. The TapeWare administrator can grant different types of permissions to various users to ensure the security and integrity of the network data while efficiently implementing a productive backup program.

In This Chapter

- Overview
- Users & Groups
- Inheriting Permissions
- Types of Permission
- Two Examples of Permissions
- Granting Permissions to Other Users

Overview

Before you can work with any object in the TapeWare storage management database, you must have **effective permissions** to that object. This is true for every object in the TapeWare database and for every user.

Different types of permissions restrict what type of operations can be performed on an object. For example, some permissions allow users to *write* to an object (such as a file, a tape, or a workstation) or *create* new objects (such as folders or jobs). Sometimes a user is granted unlimited permissions to an object or all objects. Usually, however, to protect the integrity of the data and for security reasons, most users have only limited effective permissions to some (not all) of the objects in the TapeWare database.

Maintaining the security of data on a LAN is the primary responsibility of the TapeWare administrator. Because of this, the following chapter is only an


overview of security and permissions. This chapter will help the typical user understand how TapeWare handles security permissions so that you can work efficiently with your TapeWare administrator. More detailed information on permissions is covered in Chapter 11, “Permissions and Security Reference.”

Users and Groups

For further information on setting up users and groups, see “Adding New Users and Groups,” Chapter 11.

The TapeWare administrator grants permissions to objects in the TapeWare database to either a **user** or to a **group**. Individual TapeWare users have effective permission to an object either as a user or as the member of a group. A group is a set of users that are all granted permissions in the same way and at the same time. For example, the TapeWare administrator may grant permission to read the files on a tape to users individually, to a group of users, or both to users and to groups.

Users and groups can be granted permissions to any object, in this case a tape.



Individual TapeWare users can be the member of more than one group or of every group, depending on how the TapeWare administrator arranges the storage management database security. The number of groups the TapeWare administrator creates and the assignment of members to those groups depends on the security needs of your particular LAN.

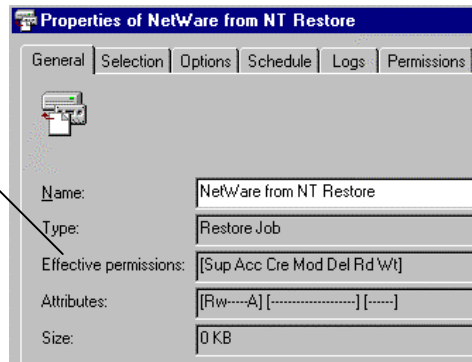
The Everyone Group

Normally most TapeWare users are a member of a special group, the **Everyone** group. Whenever a new user is added to a TapeWare storage management zone, TapeWare automatically assigns that user to the **Everyone** group. Typically, only limited permissions are granted to the **Everyone** group, although users can be granted more extensive permissions either individually or as members of other groups.

Effective Permissions

The permissions you have to an object in the TapeWare storage management database are called your **Effective Permissions**. You can view your effective permissions to an object on the **General** tab of that object.

The current user's effective permissions to an object are shown on the **General** tab of that object.



Determining Effective Permissions

A user is assigned effective permissions to an object in one of two ways, either through **direct permissions** or through **inherited permissions**.

A user has *direct permissions* to an object if they are listed on the **Permissions** tab of that object or if they are a member of a group that is listed on the **Permissions** tab of that object.

A user has *inherited permissions* to an object if (1) they do not have direct permission, and (2) they have effective permissions to the container that contains the object. This means that if you do not have direct permissions to an object, you must have effective permissions to the container in which that object is stored. (Note that your effective permissions to the container object can be either direct or inherited permissions. All that matters is that you have effective permissions to the container.)


When TapeWare determines the effective permissions a user has to an object, it first looks to see if the user has direct permission; if not, TapeWare then checks to see if the user has inherited permission.

Examples of Determining Effective Permissions

The following two examples illustrate how TapeWare determines the effective permissions a user has to an object.

Effective Permissions Example #1

In this example, a user named Galen Clark has direct permissions only to the User/Group folder named **Galen Clark's Folder**.

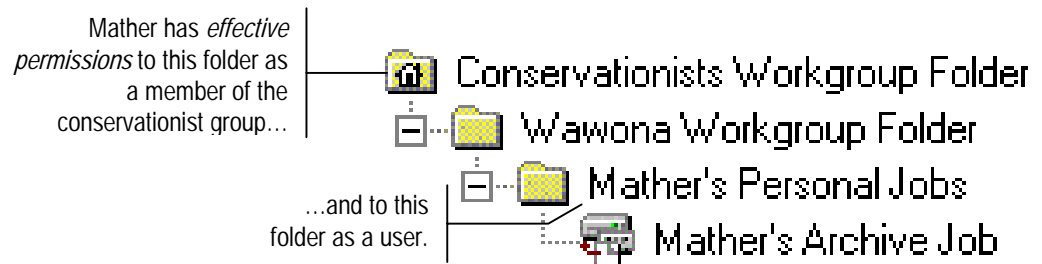


When determining the effective permissions Clark has to **Galen Clark's Folder**, TapeWare first looks to see if he has direct permission to the folder. Because Clark is listed as a user on the **Permissions** tab of the folder's property sheet, Clark has direct permission to the folder. TapeWare uses this information to determine Clark's effective rights. TapeWare does NOT look to see if there are any inherited permissions to the folder.

When determining the effective permissions Clark has to the backup job named **Clark's Personal Backup Job** stored in **Galen Clark's Folder**, TapeWare first looks to see if he has direct permission to the job. Because Clark does not have direct permissions, TapeWare checks to see if Clark has effective permissions to the container that contains the job. In this case, TapeWare checks to see if Clark has effective permissions to **Galen Clark's Folder**; because Clark has effective permission to this folder, TapeWare uses this information to calculate the effective permissions Clark has to the job.

Effective Permissions Example #2

In this example, a user named Stephen Mather has direct permission to the User/Group folder named **Conservationist Workgroup Folder** *as a member of the Conservationist Workgroup* and to the Job folder named **Mather's Personal Jobs** *as a user*.




When determining the effective permissions Mather has to the **Conservationist Workgroup Folder**, TapeWare first looks to see if he has direct permission to the folder. Because Mather is a member of the **Conservationist Workgroup** which has direct permissions to the folder, Mather also has direct permissions to the folder. TapeWare uses this information to determine Mather's effective permissions. TapeWare does NOT look to see if there are any inherited permissions to the folder.

When determining the effective permissions Mather has to the Job folder named **Wawona Workgroup Folder** stored in the **Conservationist Workgroup Folder**, TapeWare first looks to see if he has direct permission to the folder. Because no users or groups have direct permissions to this folder, TapeWare checks to see if Mather has effective permissions to the container that contains this folder. In this case, TapeWare checks to see if Mather has effective permissions to the **Conservationist Workgroup Folder**; because Mather has effective permission to this folder, TapeWare uses this information to calculate the effective permissions Mather has to the **Wawona Workgroup Folder**.

Similarly, TapeWare determines the effective permissions Mather has to the folder named **Mather's Personal Jobs** by the direct permissions Mather has to that folder. Note that when determining the effective permissions to this folder, it makes no difference that Mather also has effective permissions to the **Wawona Workgroup Folder** which contains this folder.

When determining the effective permissions Mather has to the job **Mather's Archive Job** contained in the **Mather's Personal Jobs** folder, TapeWare checks to see if Mather has direct permissions. Because he does not, TapeWare checks to see if Mather has effective permissions to the container that contains the job. Because Mather has effective permissions to the **Mather's Personal Jobs** folder, TapeWare uses these effective permissions to determine his effective permissions to **Mather's Archive Job**.

Note especially that the effective permissions Mather has to **Mather's Archive Job** are determined ONLY by Mather's effective permissions to the **Mather's Personal Jobs** folder—and NOT by his effective permissions to the **Conservationist Workgroup Folder** or the **Wawona Workgroup Folder**.




How Moving Objects Effects Permissions

When an object is moved from one container to another, TapeWare determines the effective permissions of the object based on its new location.

For example, suppose a user has effective permissions to a job because that job is stored in his or her User/Group folder, a container to which the user has been granted direct permission. If the job is moved from that folder to a new folder, the user's effective permissions to the job may change. If the job were moved to

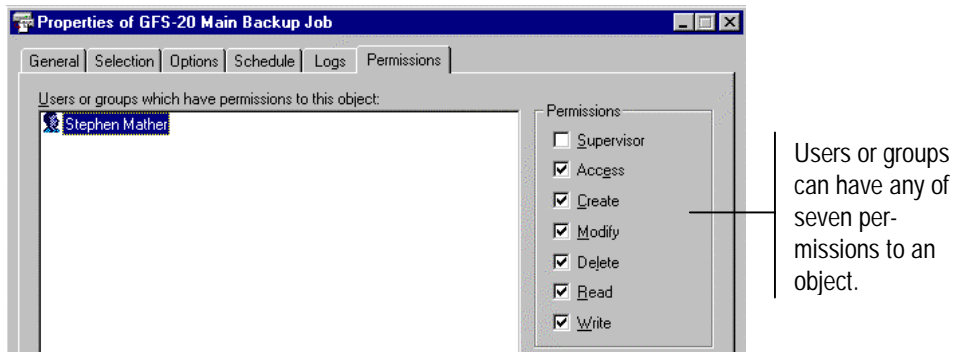
the **Admin** folder, the user would lose permission to it because he or she does not have permission to the **Admin** folder. On the other hand, if the job were moved to the **Everyone** folder, the user will still be able to access the job, although the effective permissions might be different.



Types of Permission

TapeWare controls access to objects in its database with seven different types of permission. The type of permission determines what actions a user can perform on an object. Users and groups can be granted all seven types of permission, only some of the permissions, or none of the permissions.

The seven types of permission are **Supervisor**, **Access**, **Create**, **Modify**, **Delete**, **Read**, and **Write**.



Supervisor

This is the most powerful permission. Supervisor permission grants the user three specific abilities:

- first, a user with supervisor permission to an object is *automatically granted the other six permissions to that object*;

- second, a user with supervisor permission to an object *automatically has effective permission to all the objects in the storage management database below that object*; and
- third, a user with supervisor permission to an object *cannot be denied any of the seven permissions to any object in the storage management database below that object*.

The TapeWare administrator is automatically granted supervisor permission to the highest container in the TapeWare storage management database hierarchy (called the **System Container**). This means that the TapeWare administrator has full permissions to all the objects in the TapeWare database and that none of these permissions can be denied.

Often, only the TapeWare administrator will be assigned supervisor permission.

Access

A user with **Access** permission to an object can grant other users and groups permissions to that object. For example, if you wish to grant a coworker permissions to a tape you have created, you must have **Access** permission to that tape.

Access permission can be very powerful, since it allows a single user the ability to grant all other users in the TapeWare database extensive permissions to an object. For this reason, your TapeWare administrator may not grant you **Access** permission to objects even though you have other permissions to them. For example, your TapeWare administrator may grant you permission to read and write from a particular tape. Without **Access** permission however, you will not be able to grant other users or groups that same ability.

You can only grant permissions to other users or groups if you have **Access** permission. If you want other users to have permissions to an object, such as a job or tape, and do not have **Access** permission to that object, ask your TapeWare administrator to grant the permissions for you.

Create

This permission allows a user to create new objects within a container object. For example, to create a new job within a folder, a user must have **Create** permission to the folder. Note that the **Create** permission applies *to the folder*, not to the job: it grants the user permission to create new objects *within* that folder.

If you want to create new folders or jobs, your TapeWare administrator must grant you **Create** permission. Your permission to create new jobs or folders might be limited to a single folder. For example, you may have **Create**

permission only to the **Everyone** folder or to a personal folder that your TapeWare administrator has created for you or your workgroup.

If you can't create a new job or folder, first make sure you have selected a folder in the tree view area to which you have **Create** rights. If you still cannot create a new job or folder, ask your TapeWare administrator to grant you **Create** permission to a folder.

Modify

This permission allows a user to change the name and location of an object, such as a job, in the TapeWare storage management database. **Modify** permission also allows a user to change or modify the property sheets of an object. If you have this permission, you will be able to move, rename, and change the property sheets of objects. For example, to change the name of job, a user must have **Modify** permission to that job. You must also have **Modify** permission to move a job from one folder to another.

Delete, Read, and Write

These three permissions control user access to objects, such as tapes, devices, and files, that are read, deleted, or written to. These permissions are necessary in order to run backup, restore, and verify jobs as specified below.




- *To complete a backup job*, a user must be granted **Read** permission to the files to be backed up and **Write** permission to both the media and the backup device. If the backup job will *overwrite* the media with the new data (as opposed to merely *appending* the new data), the user must also have **Delete** permission to the media.
- *To complete a restore job*, a user must be granted **Write** permission to the volumes (disk drives) on which the files are going to be restored and **Read** permission to the media and backup device. If the restore job will *overwrite* or *replace* old files, the user must have **Delete** permission to those files.
- *To complete a verify job*, a user must be granted **Read** permission to the files on the workstation or file server to be verified, to the media, and to the backup device.

Examples of Permissions

The following two examples illustrate how the permissions work with each other.

Example #1

In this example, a user named Galen Clark has effective permissions to a folder named **Clark's Jobs**, to a drive (or volume) named **Clark's 2 Gig Drive**, and to a tape named **Clark's Personal Backup Tape**.

Database Object	Effective Permissions
 <i>Clark's Jobs</i> (Job folder)	[--MCDWR]
 <i>Clark's 2 Gig Drive</i> (Volume)	[-----R]
 <i>Clark's Personal Backup Tape</i> (Media)	[-----R]

These permissions allow Clark to do the following:

- Because Clark has **Create** permission to the **Clark's Job** folder, he will be able to create new job folders within that folder and to create backup, restore, and verify jobs inside that folder. The **Modify** permission allows him to move these jobs between folders, to change the name of the job, and to change the property sheets of these jobs. The **Delete** permission allows him to delete any jobs or folders inside this folder. Clark also has **Read** and **Write** permissions to the **Clark's Job** folder and will have these same permissions to any object stored in that folder.
- Because Clark also has **Read** permission to the volume named **Clark's 2 Gig Drive**, he will be able to select files from that drive for backup. He will also be able to select files for restoring.
- Because Clark has **Read** permission to both the tape and to the volume, he will be able to create and run verify jobs—if he also has **Read** permission to the backup device.




These permissions do NOT allow Clark to do the following:

- Although Clark will be able to create a backup job, he will not be able to run the backup job because he does not have **Write** permission to the tape. The job must be run by the TapeWare administrator or some other user the TapeWare administrator grants **Write** permission to the tape.
- Similarly, although Clark will be able to create a restore job, he will not be able to run the restore job because he does not have **Write** permission to the disk drive. If Clark wants to run the restore job, he must ask the TapeWare administrator to run it for him or to grant him the permissions necessary to run it.

- Clark is prevented from granting permission to other users to his folder, to his drive, and to his tape because he lacks **Access** permission to these objects.

Example #2

In this example, a user named Stephen Mather has effective permissions to a drive named **Mather's 1 Gig Drive** and to a tape named **Mather's Personal Backup Tape**. He is also a member of the **Conservationist** group which has effective permissions to a folder named **Conservationist Workgroup Jobs**.

Database Object	Effective Permissions
 <i>Mather's 1 Gig Drive</i> (Volume)	Mather : [-----R]
 <i>Mather's Personal Backup Tape</i> (Media)	Mather : [-----WR]
 <i>Conservationist Workgroup Jobs</i> (Folder)	Conservationists : [--MCDWR]

These permissions allow Mather to do the following:

- Because he is a member of the **Conservationist** group, Mather will be able to create jobs and folders within the **Conservationist Workgroup Jobs** folder. He will also be able to modify and delete any jobs or folders in that folder.
- Because Mather has **Modify** permission to the folder and **Read** permission to the drive named **Mather's 1 Gig Drive**, he will be able to select files from that drive to backup. Similarly, Mather will be able to select files for restoring because he has **Read** permission to the tape named **Mather's Personal Backup Tape**.
- Unlike Clark, Mather will be able to run backup jobs because he has **Write** permission to the tape. (This assumes he also has **Write** permission to a tape drive.)
- Mather will also be able to create and run **Verify** jobs.

These permissions do NOT allow Mather to do the following:

- Although Mather will be able to create and run *backup* jobs, he will not be allowed to *restore* files from those tapes onto his drive because he has not been granted **Write** permission to his disk drive.
- When Mather runs a backup job, he will be prevented from overwriting old files on the tape with new files because he does not have **Delete** permissions

to the tape. Thus the backup jobs he creates and runs must all be append jobs.

- Mather, like Clark, is prevented from granting permission to other users to either his drive or his tape because he lacks **Access** permission.

Other members of the **Conservationist** group also have some permissions to the jobs and folders inside the **Conservationist Workgroup Folder**, including those created by Mather. This effects them in the following ways:

- They will be able to view the folders and jobs Mather creates in the **Conservationist Workgroup Folder**. Because they have **Modify** permission, members of the **Conservationist** group will be able to modify the properties of any jobs or folders Mather creates inside the **Conservationist Workgroup Folder**.
- Because members of this group lack **Read** permission to Mather's drive and tape, they will be unable to change the files Mather selected for backing up or restoring—even if they have **Modify** permission to a job Mather created.
- No other members of this group can run one of Mather's jobs unless the TapeWare administrator grants them the appropriate **Read** and **Write** permissions.

Granting Permissions to Other TapeWare Users

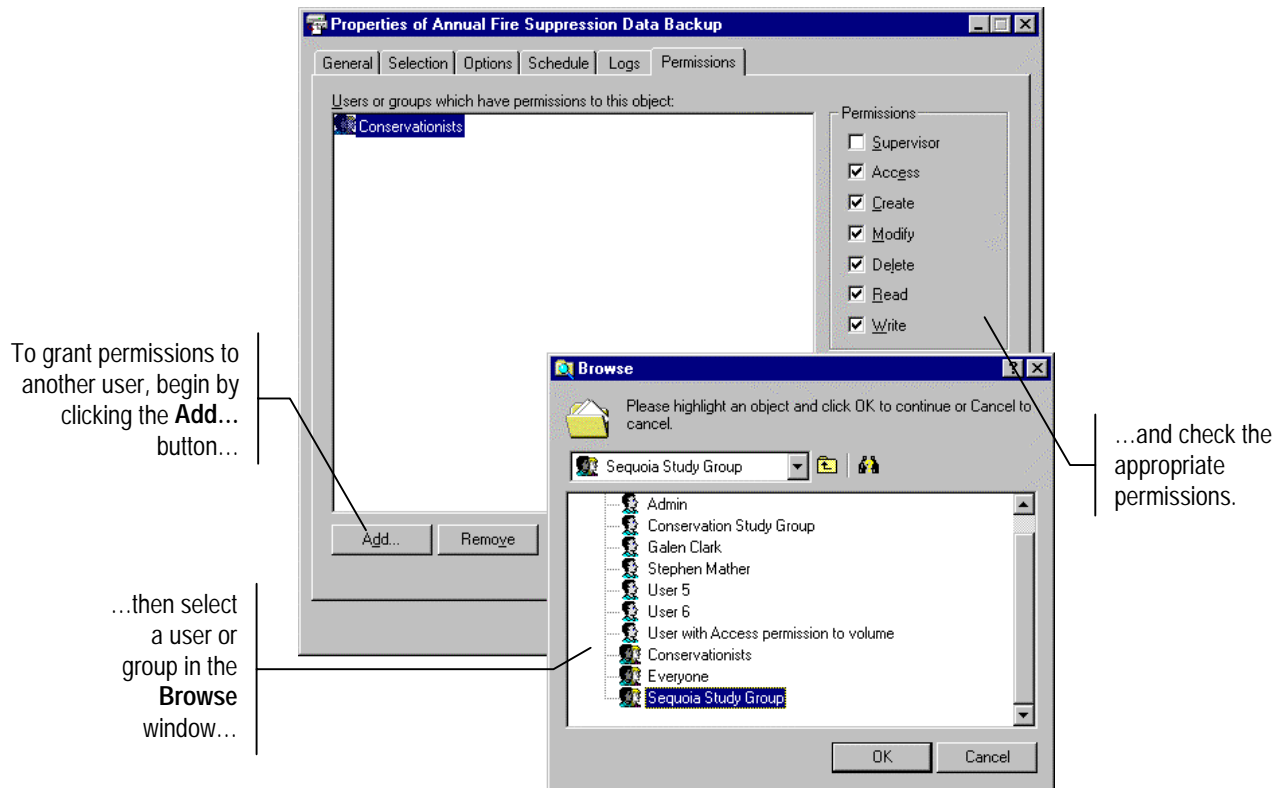
If your TapeWare administrator has given you **Access** permission to an object, you will be able to grant other users and groups permissions to that object. Generally your TapeWare administrator will only grant you **Access** permission to objects in the TapeWare database when he or she wants you to be able to share this data with other users.

For example, if your TapeWare administrator has given you **Access** permission to an archive tape, you will be able to grant other users **Read** permission to the tape. This would allow other users to restore files from this tape to any drive they have **Write** permission to.

◆ To Grant Permission to Other TapeWare Users

1. Open the properties window of the object to which you wish to grant users or groups permission.
2. Click on the **Permissions** tab. (If you don't have **Access** permission to an object, the **Permissions** tab will not appear on the object's property sheet.)
3. Click the **Add...** button.

4. Select a user or group to add in the **Browse** window and then click **OK**.
5. Select the appropriate **Permissions** check boxes at the right of the window.



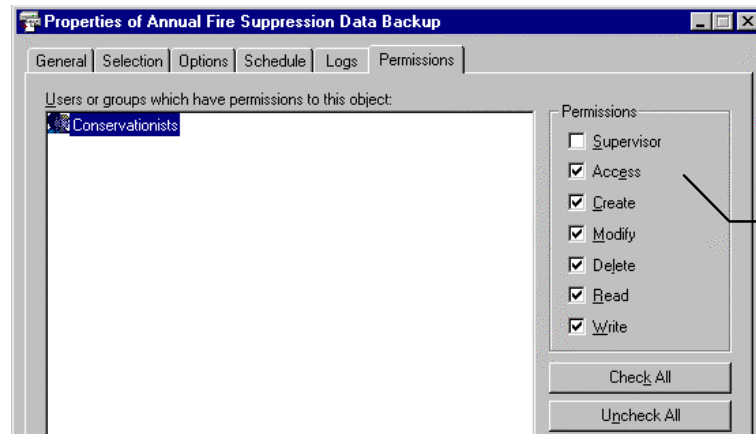
Restricting User Permissions

In most cases you can restrict the permissions a user has to an object in the same way you grant permissions—by clearing the appropriate check boxes on the **Permissions** tab of that object. If the user or group is not listed on the **Permissions** tab of the object, you must first add that user or group to the list of users or groups that have direct permissions to that object.

◆ To Restrict Permissions of Other TapeWare Users

1. Open the properties window of the object to which you wish to restrict user or group permissions.
2. Click on the **Permissions** tab. (If you don't have **Access** permission to an object, the **Permissions** tab will not appear on the object's property sheet.)

3. If the user is not listed on the **Permissions** tab, first add the user or group by clicking the **Add...** button, and then selecting a user or group to add in the **Browse** window.
4. Select the user or group to whom you wish to deny permissions on the **Permissions** tab.
5. Clear the appropriate **Permissions** check boxes at the right of the window.



To deny a user permissions to an object, add that user to the object's **Permissions** tab and clear all of the **Permissions** check boxes.

When you restrict a user's direct permissions to a container, you also change that user's effective permissions to objects within that container—but only when the user's effective permissions to those objects are *inherited* from that container's effective permissions. For example, if you deny a user direct **Modify** permission to a job folder by clearing the **Modify** check box under that user's name, you also deny that user effective **Modify** permission to jobs stored in that folder—unless that user has direct permissions to those jobs.

Normally, your TapeWare administrator will have arranged the security of your storage management zone to prevent unauthorized permission to files and tapes. However, if you believe another user's permission to an object should be restricted and you cannot restrict it yourself, notify your TapeWare administrator.

Selecting Files and Instances

You use the **Selection** tab of a job's property sheet to select files for backing up, restoring, and verifying. TapeWare's powerful selection filters allow you to select exactly the files you want and to automatically update your selection before the job is run.

In This Chapter

- Overview
- Backup Selection Concepts
- Selecting Files for Backup Jobs
- Restore Selection Concepts
- Selecting Files for Restore Jobs
- Restoring Files with New Names and Locations
- Verify Selection Concepts
- Selecting Files for Verify Jobs

Overview

When you select files for a job, you want to select only those files which are necessary for your job, and not any others. However, you also want your selection criteria to be flexible enough to automatically select new files that meet your criteria.

TapeWare allows files to be selected in several ways: by directly selecting volumes, folders, and files; by using filters to sort through selected files; and by additional automated filters used by TapeWare when your job is run. For restore jobs, you can also specify which **instance** or version of a file you wish to restore, what name that new file will have, and what folder or directory the file will be restored to.

By specifying your selection criteria carefully, when TapeWare automatically updates the selected file list each time your job is run, you can be certain that the job will work with those files you intended to select.

This chapter discusses selecting files for each of the job types separately—first backup jobs, then restore jobs, and finally verify jobs. For each type of job, there is a discussion of concepts you'll need to correctly and carefully select files. Be certain to read this section before proceeding.

Backup Selection Concepts

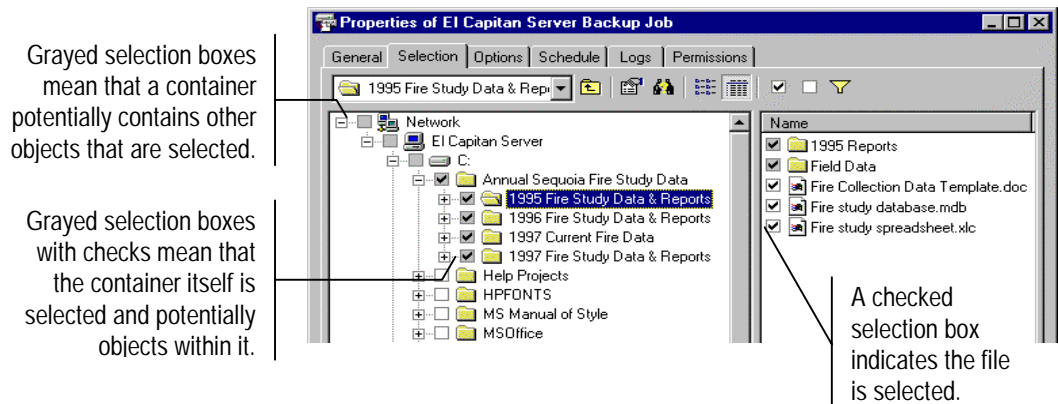
You use the **Selection** tab of a job to select files for that job.

Files are selected for backup in three steps. In the first step, the appropriate files are selected by marking them with a check. In the second step, these marked files are sorted through using file selection criteria. This step is optional. In the last step, as the job is run, TapeWare checks to see if it will backup all the files or only those files that have changed since the last backup job.

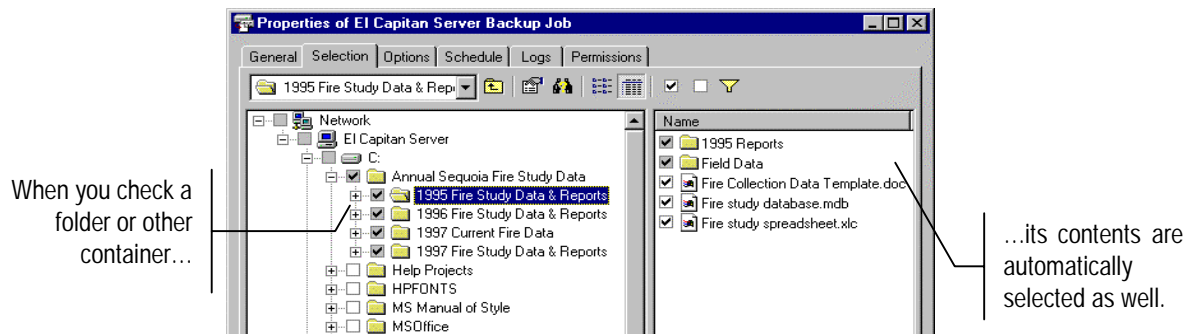
Marking Files for Backup

A file is selected for backup when a check mark appears in the selection box next to the file.

When the box next to a folder or other container is gray, that means that although this folder or container is not selected itself, it potentially contains files that are selected. When the box next to a folder or other container is *grayed and checked*, the folder itself is selected and it potentially contains files within it that are selected.




You can select or unselect a file for backup by marking or clearing the selection box next to the file. More often, however, you will want to mark the folder that contains the file, and not the file itself. Notice that when you mark a container, such as a folder or volume, all of its contents, including all the folders and containers it contains, are also marked.



Selecting Folders vs. Selecting Files

You can select the contents of the folder in one of two ways: either by individually marking the selection box of each object in that folder one by one, or by marking the selection box of the folder itself. Which method you choose is important because it effects how TapeWare recalculates the selected file list *after changes have been made to that folder*.

If you select the contents of the folder individually, when new objects such as files are added to the folder, TapeWare does not select them for backup. However, if you select the folder *itself*, when new items are created in that folder, TapeWare also selects these files for backing up.

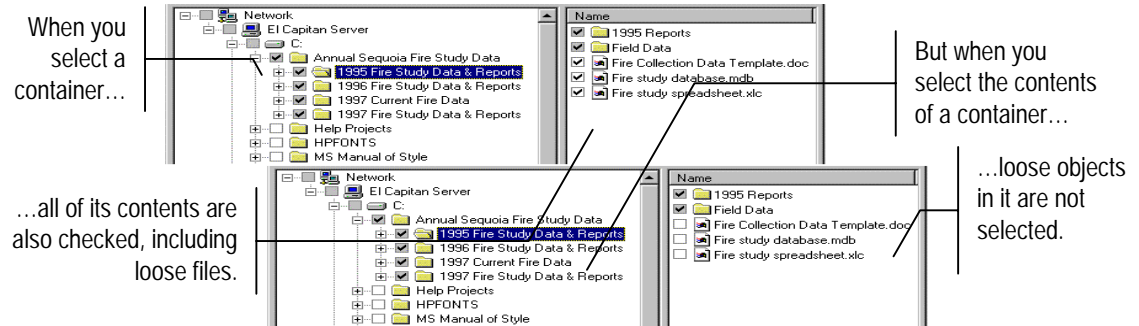


In general, when selecting files for backup, especially for jobs designed for disaster protection, begin by selecting containers at the top of the TapeWare hierarchy. Then, as needed, unselect containers or files lower in the storage management database hierarchy that you do not require.

For example, you could begin by selecting the network container icon at the top of the hierarchy. This will automatically select all of the machines on the network and all of the volumes on those machines. If there are machines, volumes, or folders you do not want backed up, unselect them by clearing their selection box. When new machines or volumes are added to the network (that is to the current storage management zone), TapeWare will automatically select these machines and volumes.

There is an additional reason to select containers rather than the objects in the containers: to insure that you don't miss any files stored loose in the container.

For example, suppose you wanted to backup a folder named **Workgroup Memos**, including the folder stored in that folder named **Confidential Memos**, but not the folder named **Party Announcements**. If you just checked the **Confidential Memos** folder, your job would miss any files stored 'loose' in the **Workgroup Memos** folder. Instead, begin by checking the **Workgroup Memos** folder and then clearing the **Party Announcements** folder. This way you can be certain to select any files stored loose in the **Workgroup Memos** folder.



Sorting Files with Filters



Selection Filters
button

You can also sort files for backup with filters by clicking on the **Selection Filters** button on the toolbar and specifying filters which exclude files that do not meet the filter criteria. This is an optional step.

The selection filters you apply are applied to all of the volumes, folders, and files that have been marked for backup. *You cannot apply different filters to different folders or volumes.*

Filters Exclude, not Include

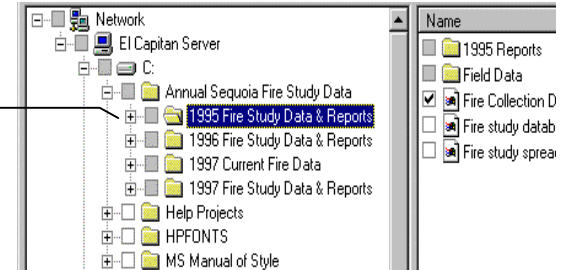
The selection filters exclude files by filtering out files that do not meet the selection criteria. If a folder or other container has been marked for backup, TapeWare uses the selection filters to sort through the files and unmark any files that do not meet the selection criteria. *TapeWare does not use the selection filters to add files to the backup set.*

The filter criteria are applied to all marked files, regardless if they were marked before or after the filter criteria were specified. After specifying the selection filter criteria, you can then mark or unmark files, folders, and volumes for backup. TapeWare will apply the filter criteria to whatever files and folders are selected whether they are marked before or after the filter criteria are specified. Additionally, you can change the filter criteria at any time; TapeWare will automatically reapply the new selection filter criteria to the marked folders and files.

Filters and Grayed Folders

When you apply a filter to a marked folder, this may result in no files from that folder being selected. Although no files in that folder are selected, the folder will still appear with a grayed selection box next to it. The grayed selection box indicates that if any new files which meet the filter criteria are created in that folder, they will be selected for backup.

Grayed folders indicate that any files contained in the folder that meet the selection filter's criteria will be selected.



Selecting Changed Files only

When you run a job repeatedly, particularly backup jobs designed for disaster protection, many times you only want to back up files that have changed since the last time you've run a backup job. Normally, TapeWare automatically handles this step.

When TapeWare runs a job, it checks to see whether you want the job to back up all selected files or only those files that have changed. If you want to back up only files that have changed since the last backup, TapeWare automatically unselects all the files that have not changed. The job then runs with this updated selected file list.


Automatically Selecting New Files for Backup

When you set up a job to run repeatedly, you want that job to adjust to changes made to the directories on the LAN. Sometimes these changes include the addition of new files and directories that were not originally selected when the job was created. It may even include the addition of new machines on the LAN or volumes on these machines that were not previously selected for backup.

If your selection criteria was carefully specified, TapeWare will also select these new files, folders, volumes, and machines for backup. In general a new file or container is selected for backup if (1) it is contained in a container selected for backup and (2) it passes the selection filter criteria.

For example, a new folder will be included in the backup selection list if it is contained in a container that was itself selected. If you create a new folder at the root level on a volume, that folder will be selected for backup if the volume was

marked for backup. The files in that folder that pass through the selection filter criteria will be included in the backup job. However, if the volume itself was not selected, the new folder will not be selected.



If you want to see what files will be backed up, open up the **Selection** tab on the job's property sheet. Whenever this tab is opened, TapeWare recalculates which files should be selected for backup. Check the display to see that the files you intended to be backed up are selected.

Selecting Files for Backup Jobs

◆ To select files for backup jobs

1. Open the property sheet of the backup job and click on the **Selection** tab.
2. Mark the selection boxes next to the folders or containers you wish to include in the job.
3. Click on the **Selection Filters** button on the Toolbar and specify a filter selection criteria. (You can skip this step if you don't wish to apply any selection filters.)
4. Check the tree view area and object detail area on the **Selection** tab to see that the files you intended to select are marked for backup.

Applying Filter Criteria

When you click on the **Selection Filters** button, the **Selection Filters** window appears. This window has multiple selection filters you can apply to the files you have selected.

The Backup Job
Selection Filters
window.

Note that each filter criterion works independent of every other. In order to be selected for backup, each file must pass every filter criterion specified. For example, if you specify that every file selected must have been created after January 1, 1997 and must have .doc as its extension, TapeWare will only select files which meet *both* selection criteria.

The Filter Selection Criteria

This section contains a brief description of each selection filter TapeWare applies to the files and folders marked for backup.

Note that some of the selection criteria are operating system specific. Your storage management zone may include multiple machines working with files created by different operating systems. If you select a filter criterion that is operating system specific, files from other operating systems will be automatically excluded from the backup. This effects in particular the **Required Attributes** and **Exclude Attributes** filters.

Note TapeWare figures the century dates using the following algorithm: if the year date is 70 or smaller, the century date is set to 20 (21st century); if the year date is 71 or greater, the century date is set to 19 (20th century). For example, if you set the date to 6/1/33, TapeWare calculates the date as June 1, 2033. If the dates is set to 4/5/81, TapeWare calculates the date as April 5, 1981.



Selection Filters
button

Backup Range

The **backup date** is the date a file or other object was last backed up. Each time a file is backed up, TapeWare assigns changes the backup date to match the current date. Each file has exactly one backup date. You can use this filter files that have been backed up on specific dates. More often however you would use this filter to filter out files that have been recently backed up



Filter Change
button

To select files that have specific backup dates, click on the **Filter Change** button next to the **Backup range** field. In the **Date Range** window that appears, select the appropriate range type and the starting and ending dates and times.

To select files by
backup date, set the
Backup Range filter to
the desired date.

The **Date Range** dialog box has a title bar with a question mark and close button. It contains a 'Range type:' dropdown menu set to 'Not between two dates'. Below this are four input fields: 'Starting date:' (01-Jan-70), 'Starting time:' (12:00 Am), 'Ending date:' (06-Feb-06), and 'Ending time:' (06:28 Am). At the bottom are 'OK' and 'Cancel' buttons.

Modify Range

Each time a file is modified, its modified date is updated. You can use this filter to backup files whose modified date matches your criteria. TapeWare checks the directory information on the volume to see if the file should be included for backup. For example, you can select only those files that were modified *after* a certain date and time or, alternatively, those that were modified *before* a certain date and time.



Filter Change
button

To specify a range of dates, click on the **Filter Change** button next to the **Modify range** field. In the **Date Range** window that appears, select the appropriate range type and the starting and ending dates and times.

Date Range
window

The **Date Range** dialog box has a title bar with a question mark and close button. It contains a 'Range type:' dropdown menu set to 'Between two dates'. Below this are four input fields: 'Starting date:' (01/01/1970), 'Starting time:' (12:00 Am), 'Ending date:' (02/06/2106), and 'Ending time:' (06:28 Am). At the bottom are 'OK' and 'Cancel' buttons.

Create Range

When a file is first created, it is assigned a create date. You can use this filter to select only those files which match your criteria. TapeWare checks the created date for each file stored in the directory of the volume and uses this to select files for backing up.



Filter Change
button

For example, you might want to back up only those files created after a certain date. To do so, click on the **Filter Change** button next to the **Create Range** field and then select the appropriate criteria in the **Date Range** window that appears.

Access Range

Each time a file is read, whether or not it is modified, its access date is updated. You can use this information to select files for backup. For example, you might want to backup only those files which have been accessed (opened or read) in the past two months. To do so, in the **Date Range** window, select **On or after** in the **Range type** field. Then indicate the appropriate starting date and time.

Alternatively, you could back up only those files which have *not* been accessed in the past two months by selecting **On or before** in the **Range type** field.

Note If the operating system does not support access dates or create dates for files, this filter is ignored. For example, DOS does not support access dates or create dates. If you specify a filter for either one of these dates, this filter will be ignored when TapeWare encounters files on a DOS volume.

Size Range



Filter Change
button

This filter lets you select files for backup according to their size. You might want to select only smaller files, larger files, or files between two sizes. To specify a filter that sorts files according to their sizes, click on the **Filter Change** button next to the **Size range** field and then select the appropriate criteria in the **Size Range** window that appears.

Instance Range

Each time TapeWare backs up a file, it creates a new **instance** of that file. For example, a file named **Expense Account Reporting Form** may have been backed up several times during the previous months and years. Typically, each instance of the file is stored on the backup media of a different job. TapeWare tracks each instance of a file separately in its storage management database.

You can use this filter to instruct TapeWare to not back up files for which you already have multiple instances. For example, you may not want any more than three instances of a particular file backed up. Set the **Range Type** to **At most**

and then set the **Maximum instances** field to **3**. TapeWare only backs up those files for which there exist less than 3 instances.

Note however that having multiple instances of a file does not insure that the instances you have reflect the latest changes to the file. It may have been modified after the last time you backed it up and so your latest instance may not match the file's current form.

Must Match

TapeWare lets you sort files using wildcard matches. Only files that match the wildcard indicated in the **Must match** field are included in the backup set. For example, if you enter “*.exe” TapeWare will only back up those files with the .exe file extension.

You can specify multiple wildcards by separating each with a semicolon, “;”. For example, if you enter “*.exe;*.doc” in the **Must match** field, TapeWare selects all files that *either* have the .exe extension *or* the .doc extension.

Wildcards can be in any one of three formats: DOS, Long, or UNIX. Consult your operating system documentation for more information about each wildcard format.

Cannot Match

This wildcard field works just like the **Must match** field except that it *excludes* any files that match the wildcards. You can specify multiple wildcards by separating them with a semicolon; if you specify multiple wildcards, TapeWare excludes any file that matches any one of the wildcards you specify.

Wildcard Type

You can use one of three types of wildcard formats: DOS, Long, or UNIX. Select which wildcard format you wish to use from the list box.

Required Attributes

Operating systems track certain features of files called *attributes* that they use to manage these files. You can use these same attributes as a selection filter. In the **Required Attributes** field, if an attribute is checked, TapeWare only selects those files which have these attributes. For example, if you check **Hidden**, TapeWare only selects for backup those files which the operating system has assigned the **Hidden** attribute.

You can select multiple attributes. In this case, TapeWare only selects those files that meet *all* of the required attributes.

Note that some of these attributes are only supported by certain operating systems. If you specify an attribute that is specific to a particular operating system, then only files created under that operating system will be selected for backup.

Exclude Attributes

This field works like the **Required Attributes** field except that TapeWare excludes files that match these attributes. For example, if you have checked the **Execute Only** box, TapeWare will exclude from the backup any files with the **Execute Only** attribute.

You can select multiple attributes. If you do, TapeWare excludes any files that has *any one* of the attributes. For example, if you select the **Hidden** and **System** attributes, a file will be excluded if it has *either* the **Hidden** attribute *or* the **System** attribute.

Parents

When this option is checked, when TapeWare backs up a file, it also backs up the directory information for the parent of that file. This option must be checked in order for folders and other directory data to be backed up. When this option is not checked, TapeWare will not back up any parent information for any file backed up. If unchecked, directory information about folders and volumes is not backed up.

Children

When this option is checked, TapeWare backs up the selected files. If you want only to back up the marked directories, however, you can clear this option. When the **Children** box is unchecked and the **Parents** box is checked, TapeWare backs up the directory structure, but not the files stored in the directories (that is, in the folders).

This option can be useful for replicating a complex directory structure. Begin by marking the directory structure you wish to duplicate. Then clear the **Children** option. TapeWare will backup only the directory structure. You can then replicate that directory structure to any volume by restoring the directory to that volume.

Media

TapeWare tracks instances of files and the media on which those instances are stored. You can use this information to sort files according to the media on which they appear. Only files with instances on the media in the **Media** field will be selected for the backup job. For example, if you select media named "Daily Set:1," TapeWare will only include files in the backup job which have a valid instance on the on the media named "Daily Set:1."

To sort files according to the media on which they appear, click the **Add...** button and select the media from the **Browse** window. Note that you must select a Media object, not a Media Folder or User/Group Folder. If there are multiple media shown in the **Media** filter field, only files which have a valid instance on *all* the media listed will be selected.

This filter has limited applications for backup jobs. One way to use it, however, would be in the case of media you know is corrupt or damaged. To back up a new instance of every file on that damaged media, begin by creating a new job and then selecting the appropriate Network object or Machine objects on the **Selection** tab. Then open the **Selection Filters** window by clicking on the **Filters** button. Add the damaged media to the **Media** field. Then set the **Backup mode** on the job's options tab to **Snapshot**. When TapeWare runs the job, it will only backup files which had an instance on the damaged media shown in the **Media** field.

Restore Selection Concepts

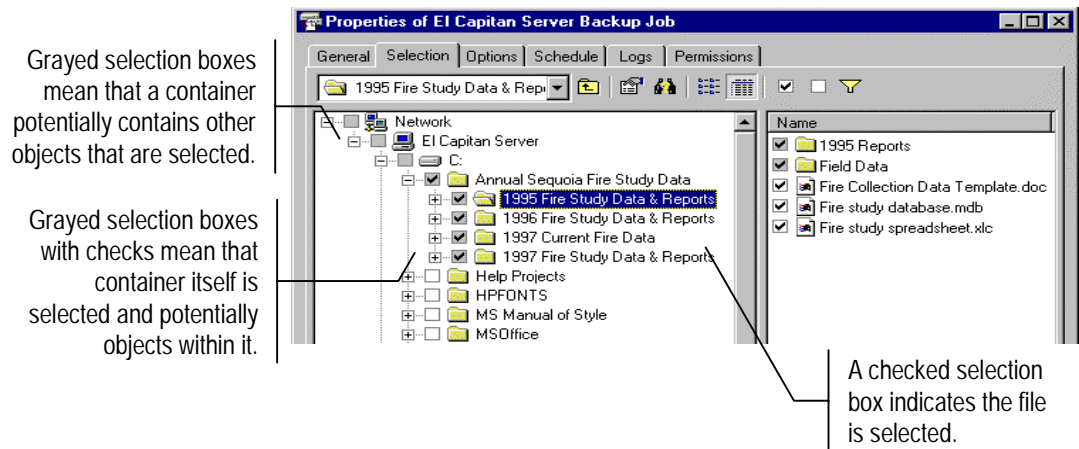
You select files for restoring the same way you select files for backing up; however besides selecting which files you wish to restore, you can also change the name of the restored file, restore it to a new location, and create a new folder in which to restore the file. Additionally, when you select a file for restoring, you must specify which instance of the file you wish to restore.

Files are selected for restoring in four steps. In the first step, you modify the file tree to look the way you want it to appear when you restore the files. For example, you might make a new folder to store the restored files in. Second, the appropriate files are selected by marking them with a check and selecting the appropriate instance. In the third step, these files are filtered using multiple selection criteria. In the last step, you can specify new names and locations for the restored files.

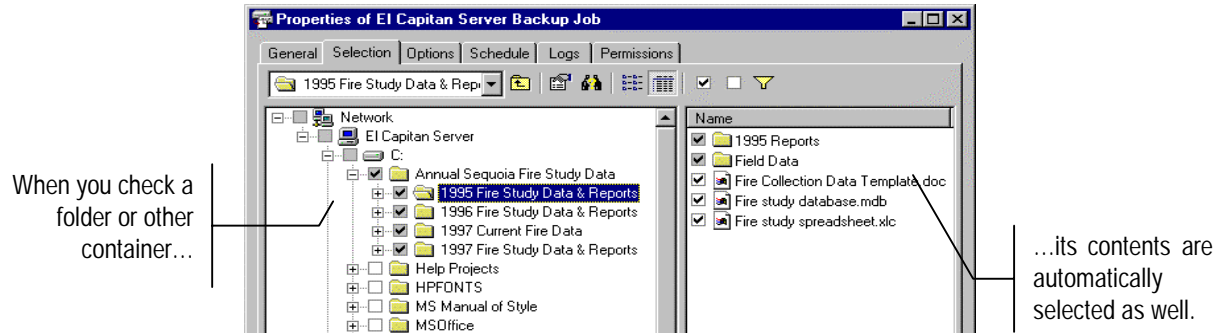
Selecting Files for Restoring

A file is selected for restoring when a check mark appears in the selection box next to the file.

When the box next to a folder or other container is gray, that means that although this folder or container is not selected itself, it potentially contains files that are selected. When the box next to a folder or other container is *grayed and checked*, the folder itself is selected and it potentially contains files within it that are selected.



You can select or unselect a file for restoring by marking or clearing the selection box next to the file. You can also mark the folder that contains the file, and not the file itself. Notice that when you mark a container, such as a folder or volume, all of its contents, including all the folders and containers it contains, are also marked.




Note that the tree view area and object detail area are different for restore jobs than for backup jobs. For restore jobs, the files displayed in the tree view area and object detail area are the files for which TapeWare has instances in its storage management database. For backup jobs, on the other hand, the files displayed are those currently present on the file servers and workstations in the current database zone.

Selecting Instances of Files

Each time a file is backed up an *instance* of that file is created. There may be multiple instances of files stored on different media created by different backup jobs. TapeWare keeps track of all the instances of each file in its storage

management database and the media on which each instance is stored. When media is overwritten or deleted, TapeWare deletes those instances from its storage management database as well.

When you select a file for restoring, TapeWare automatically selects the **<Latest>** instance. If you want to select an instance other than the latest instance of a file backed up, you must select that instance in the **Instances** window. When you open the **Instances** window, the **Available instances** field shows a list of the instances of the file and the media on which those instances are stored. Select which instance you want to restore by highlighting it. For more information about a particular instance of a file, click on the **Details** button. TapeWare displays various details it uses to manage the file in its storage management database, including its backup date and its modify date.



If you select the **<Latest>** parameter, TapeWare will restore the most recent instance of that file or folder.

Note that the files included in a restore job need not all have the same instance date. You can specify which particular instance of a file you would like to have restored for each and every file. The instance date of any particular file can be specified individually for that file.

Instances and Filters

You cannot use filters for selecting instances. The **Selection Filters** window can be used to sort through the instances you have otherwise specified in the **Instances** window, but filters will not change the instance date selected.

Selecting Instances and Folders

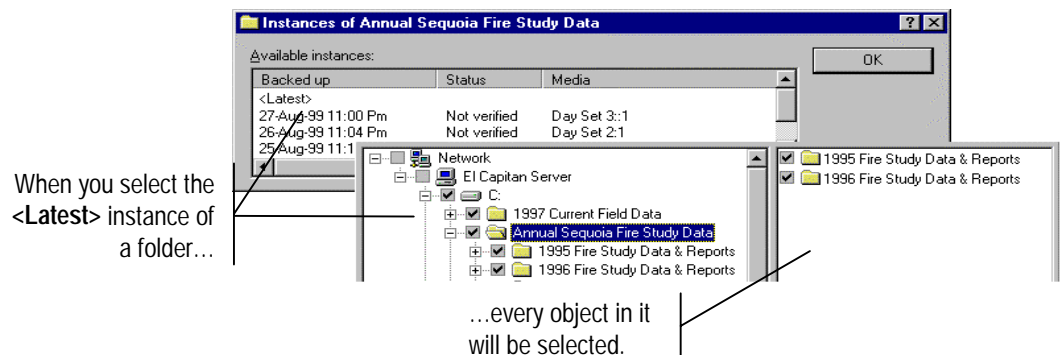
When you select a folder, TapeWare automatically selects the **<Latest>** instance for that folder and for every file within that folder. If you wish to specify another instance, open the **Instance** window and highlight the instance date to select it.

What instance you specify when you select the folder is also used to select files contained within that folder. Specifically, a file is selected for restoring only if it has an instance that matches the instance of folder.

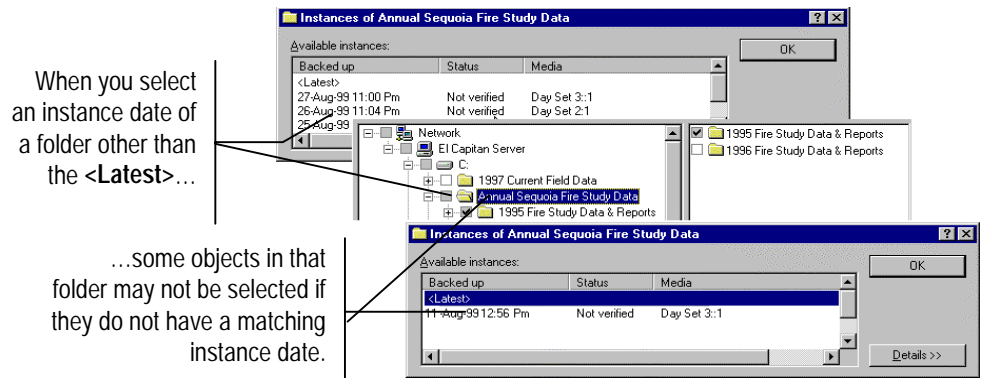
Note When you specify an instance date for a folder, volume, or other container, files stored in that container are *only selected when they have an instance date that matches the instance date of the container*. Many times, files will not have instance dates that match the dates of the containers they are stored in, for example, when you select an instance date from an incremental or differential backup job. If you want to be certain to select all of the files inside a container, select the **<Latest>** instance for that container.

Consider these two examples.

In the first example, by selecting the **<Latest>** instance of the folder, all of the files contained within the folder are selected because each of these files has a **<Latest>** instance. Note that *these instances may be from different dates and different backup jobs*, but because each file has a **<Latest>** instance, each one will be restored.



In the second example, another instance besides the **<Latest>** instance was selected. In order for the files within the folder to be selected for restoring, they must have an instance date that matches that selected for the folder. In this example, some files are not selected for restoring because they do not have an available instance which matches the instance date selected for the folder.



In general, if you want to restore a specific instance of the file, you must select that file directly and specify which instance you wish to restore in the **Instances...** window.

Selecting Folders vs. Selecting Files

You can select the contents of the folder in one of two ways: either by individually marking the selection box of each object in that folder one by one, or by marking the selection box of the folder itself. Which method you choose is important because it effects which files TapeWare includes in the selection list *after changes have been made to that folder*.

For example, if you select a folder for restoring by marking its selection box, all of the contents of that folder are restored. If a new backup job is run before the restore job is run, TapeWare selects files for restoring using the new folder's contents. So for example, if a new file is created in that folder, TapeWare will also restore that file. Additionally, if you have selected the **<Latest>** instance of the folder, TapeWare will use the latest instance of each file in its storage management database. These files may be newer than the files you originally selected.

Selecting Files and Instances with Filters

You can also sort files for restoring with filters by clicking on the **Selection Filters** button on the toolbar and specifying which file types to include or exclude.

The selection filters you specify are applied to all of the volumes, folders, and files that have been marked for restoring. *You cannot apply different filters to different folders or volumes.*

Filters Exclude, not Include

The selection filters exclude files by filtering out files that do not meet the selection criteria. If a folder or other container has been marked for restoring, TapeWare uses the selection filters to sort through the files and to unmark any files that do not meet the selection criteria. *TapeWare does not use the selection filters to add files to the restore set.*

Changing the Name and Location of Restored Files

When you restore a file, you may wish to restore the file with a new name or in a new location. If you restore a file to its original location using its original name, if that file currently exists there, TapeWare overwrites the current file with the restored file.

You can avoid overwriting current files by giving the file a new name or by restoring the file to a new directory. For example, to avoid replacing the current file named **Project List** with a previous, older instance of the file, you can rename the file before restoring it or restore it to a different folder.

You can either select a different folder, or, alternatively, TapeWare lets you create a new folder in which to restore the files.

Selecting Instances of Files for Restore Jobs

◆ To select instances of files for restore jobs

1. Open the property sheet of the restore job and click on the **Selection** tab.
2. Check the selection boxes next to the files, folders or other containers you wish to include in the job.
3. To select a specific instance of the objects you selected, highlight the folder or file and click the **Select Instance** button. In the **Instance** window that appears, select the appropriate date of the instance you wish to restore. If you want to include all of the files in a folder or on a volume, select **<Latest>**.
4. Click on the **Selection Filters** button on the Toolbar and specify a filter selection criteria. (You can skip this step if you don't wish to apply any selection filters.)
5. Examine the tree view area and object detail area on the **Selection** tab to check that the files you intended are marked for restoring.



The Select Instance button

Additionally, you can change the names of the files and store them in new locations. This is discussed in the next section, “Restoring Files with New Names and Locations.”

Selecting Specific Instances

Each time you mark a selection box of a file or folder for restoring, TapeWare automatically selects the **<Latest>** instance. If you want to specify a different instance, use the **Instances...** window to select which instance of the file you wish to restore.



The Select Instance button

You can also specify a particular instance of a file or folder by highlighting it in the tree view area or object detail area and then clicking the **Select Instance** button on the toolbar. TapeWare will show you the **Instances...** window with a list of the available instances for that file.

Be certain to specify carefully which instance of a file you wish to restore. A single TapeWare restore job can restore files backed up over a period of months or years on media created by many different backup jobs. You can easily restore all of the latest instances of the files by selecting **<Latest>** in the **Instances...** window. However, if you want instances of files that were created on different dates, you must select each instance of each file individually.

Note that you cannot restore multiple instances of the same file in a single restore job. If you want to restore more than one instance of a file, you must create and run a separate job for each instance.

Applying Filter Criteria

When you click on the **Selection Filters** button, the **Selection Filters** window appears. This window has multiple selection filters you can use to sort through the files you have selected for restoring.

The Restore
Job Selection
Filter window

Note that each filter criterion works independent of every other. In order to be selected for restoring, each file must pass every filter criterion specified. For example, if you specify that every file selected for restoring must have been created after January 1, 1997 and must have .doc as its extension, TapeWare will only select files which meet *both* selection criteria.

The Filter Selection Criteria

This section contains a brief description of each selection filter TapeWare applies to the files and folders marked for restoring.

Note that the **Selection Filters** window for restore jobs is similar to the **Selection Filters** window for backup jobs. This allows you to use the same filters to select the files for restoring that you used for selecting files for backing up previously. This allows you to create a restore job that selects the same files as a backup job, no matter how widely distributed over the network these files may be.

Backup Range

When a file is backed up, TapeWare stores in its storage management database the date the file was backed up. This is called the backup date. Each time you back up a file, TapeWare changes the backup date to the date of the backup. (You can view this information for all of the available instances in the **Instances...** window in the **Backed up** field list.) You can use this information to filter files for restore jobs.

Modify Range

Each time a file is modified, its modified date is updated. You can use this filter to restore files with a modify date that matches your criteria. TapeWare checks the directory information on the volume to see if the file should be included in the restore job. For example, you can select those files that were modified after a certain date and time or, alternatively, those that were modified before a certain date and time.



Filter Change
button

To specify a range of dates, click on the **Filter Change** button next to the **Modify range** field. In the **Date Range** window that appears, select the appropriate range type and the starting and ending dates and times.

The screenshot shows a 'Date Range' dialog box with the following fields and values:

Field	Value
Range type	Between two dates
Starting date	01-Jan-70
Starting time	12:00 Am
Ending date	06-Feb-06
Ending time	06:28 Am

Buttons: OK, Cancel

Date Range
window

Create Range

When a file is first created, it is assigned a create date. You can use this filter to select only those files which match your criteria. TapeWare checks the created date for each file stored in the directory of the volume and uses this to select files for restoring.



Filter Change
button

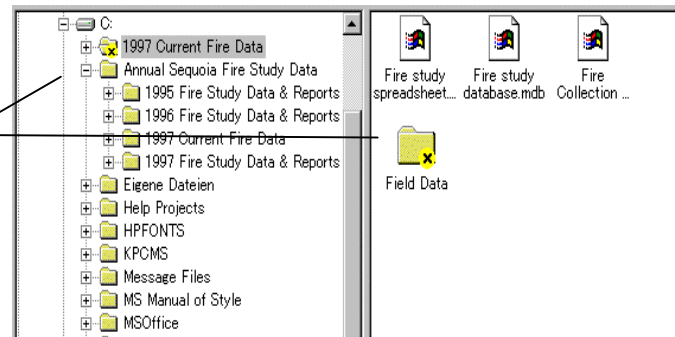
For example, you might want to restore only those files created after a certain date. To do so, click on the **Filter Change** button next to the **Create range** field and then select the appropriate criteria in the **Date Range** window that appears.

Delete Range


This filter gives you an easy way to select files for restoring that have been deleted from the volume, but for which TapeWare has valid instances stored in its storage management database and on valid media.

When files that have once been backed up are later deleted, TapeWare marks these files with a special icon, indicating that they have been deleted. In addition, TapeWare assigns the file a delete date, which you can view on the **General** tab of that file's property sheet.

Files that have been backed up and *then deleted* show special icons on the **Selection** tab of a Restore job.



Any object that has been backed up and then is deleted will show a **Delete** date on its **General** tab.



This filter instructs TapeWare to only restore files which have a delete date that matches the criteria you have set. Note that if a file has not been deleted from the volume, it will be excluded by this filter and thus will not be selected for restoring.

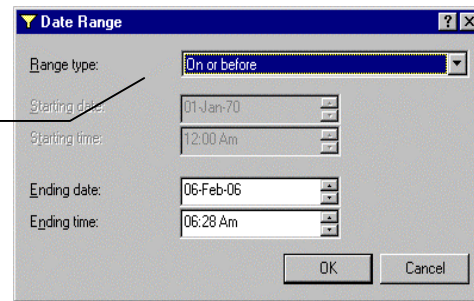


Filter Change button

You can easily select every file deleted for restoring by first marking the selection box of the volume and selecting the **<Latest>** instance. This will cause all of the files to be initially selected. Then, click on the **Filter Change** button next to the **Delete range** field and then select **On or before** from the **Range type** list box. Next select a random future date, for example, February 6, 2106. TapeWare will exclude all of the files that have not been deleted from the set of files to be restored. When you return to the **Selection** tab, all and only those files that have been deleted will be checked.

On the other hand, you can also *not* restore files that have been deleted. In this case, set the **Delete range** filter to **On or after** some random early date, such as January 1, 1980. Any file that has been deleted will be filtered out by this filter, so that no deleted files will be restored. This can be useful if you don't want to unnecessarily restore files that were properly deleted in the first place.

To select every file that has been deleted for restoring, set the **Delete range** filter to **On or before** some random future date.



Access Range

Each time a file is read, whether or not it is modified, its access date is updated. You can use this information to select files for restoring. For example, you might want to restore only those files which have been accessed (opened or read) in the past two months. To do so, in the **Date Range** window, select **On or after** in the **Range type** field. Then indicate the appropriate starting date and time.

Alternatively, you could restore only those files which have *not* been accessed in the past two months by selecting **On or before** in the **Range type** field.



Filter Change
button

Size Range

This filter lets you select files for restoring according to their size. You might want to select only smaller files, larger files, or files between two sizes. To specify a filter that sorts files according to their sizes, click on the **Filter Change** button next to the **Size range** field and then select the appropriate criteria in the **Size Range** window that appears.

Instance Range

Each time TapeWare backs up a file, it creates a new instance of that file. For example, it might have backed up a file named **Expense Account Reporting Form** several times during the previous months and years. Typically, each instance of the file is stored on the backup media of a different job. TapeWare tracks each instance of a file separately in its storage management database.

You can use this filter to instruct TapeWare to select files according to the number of instances that exist in the database. You might, for example, instruct TapeWare to restore all of the files for which there is only one instance. When you set the **Instance range** filter to **At most 1**, TapeWare only restores those files which have a single instance.

Must Match

TapeWare lets you include files using wildcard matches. If the file matches the wildcard indicated in the **Must match** field, TapeWare includes it in the restore

set. For example, if you enter “*.exe” TapeWare will only restore those files with the .exe file extension.

You can specify multiple wildcards by separating each with a semicolon, “;”. For example, if you enter “*.exe;*.doc” in the **Must match** field, TapeWare selects all files that *either* have the .exe extension *or* the .doc extension.

Cannot Match

This wildcard field works just like the **Must match** field except that it *excludes* any files that match the wildcards. You can specify multiple wildcards by separating them with a semicolon; if you specify multiple wildcards, TapeWare excludes any file that matches any one of the wildcards you specify.

Wildcard Type

You can use one of three types of wildcard formats: DOS, Long, or UNIX. Select which wildcard format you wish to use from the list box.

Required Attributes

Operating systems track certain features of files called attributes that they use to manage these files. You can use these same attributes as a selection filter. In the **Required Attributes field**, if an attribute is checked, TapeWare only selects those files which have these attributes. For example, if you check **Hidden**, TapeWare only selects those files which the operating system has assigned the **Hidden** attribute.

You can select multiple attributes. In this case, only files which have all of the specified attributes will be selected.

Exclude Attributes

This field works like the **Required attributes** field except that TapeWare excludes files that match these attributes. For example, if you have checked the **Execute Only** box, TapeWare will exclude from the verify job any files with the **Execute Only** attribute.

You can select multiple attributes. A file that has any one of these attributes will be excluded. For example, if you mark the **Hidden** and **System** attributes, any file that has *either* attribute will be excluded.

Parents

When this option is checked, TapeWare restores directory information for any selected folder or volume. For example, if you have marked a folder, TapeWare will restore that folder only if this option is checked. When this option is not checked, directory information about folders and volumes is not restored.

Children

When this option is checked, TapeWare restores files. When this option is unchecked, TapeWare does not restore files. This is useful if you want to restore a complex directory structure, but not the files in that directory. To restore a directory structure but not the files (children) stored in the directories, begin by marking the directory for restoring. Then clear the **Children** option. TapeWare will restore only the directory structure to the volume you specify.

Media

TapeWare tracks instances of files and the media on which those instances are stored. You can use this information to sort files according to the media on which they appear. Only files with instances on the media in the **Media** field will be selected for the restore job. For example, if you select media named “Daily Set:1,” TapeWare will only include files in the restore job which have a valid instance on the on the media named “Daily Set:1.”

To sort files according to the media on which they appear, click the **Add...** button and select the media from the **Browse** window. Note that you must select a Media object, not a Media Folder or User/Group Folder. If there are multiple media shown in the **Media** filter field, only files which have a valid instance on *all* the media listed will be selected.

For more information on importing media, see “Import Media,” Chapter 9.

This filter can be useful for restore jobs if you want to restore files only from a particular media. For example, you may have imported media from another storage management zone and may wish to limit the files restored to those on that particular media. In this case, you can insure that only files on that media are selected by adding that media to the **Media** field.

Note, however, that under ordinary circumstances, you should let TapeWare track the instances of particular files and restore files *not according to the media on which they appear*, but rather *according to their instance date*. For example, if you want to restore the most recent version of a file, simply select the file. TapeWare will automatically select that file and identify the proper media on which that instance is stored. TapeWare will then prompt you for the correct media when the job is run.

Restoring Files with New Names and Locations


You can change the name and location (folder) of a file when it is restored. You can also create a new folder in which to store the file.

Restoring a File with a New Name

After a file has been selected for restoring, you can rename the file. When you rename the file, TapeWare restores the file with the new name. This can be useful for not overwriting versions of the file that currently exist on disk.

To rename a file, highlight it and select **Rename** from the **Shortcut** menu. You can also rename a file by selecting its name again after it has been highlighted.

To restore a file with a new name, highlight it and then select it again. Then type the new name.



Note that when you rename an instance, you are *only* renaming that file for the purposes of restoring it with this particular restore job. *Only the current restore job will assign that file the new name.* When you create a new restore job, you will see the file displayed with its original name. Similarly, the **Database** tab always displays files with the names they had when they were backed up.

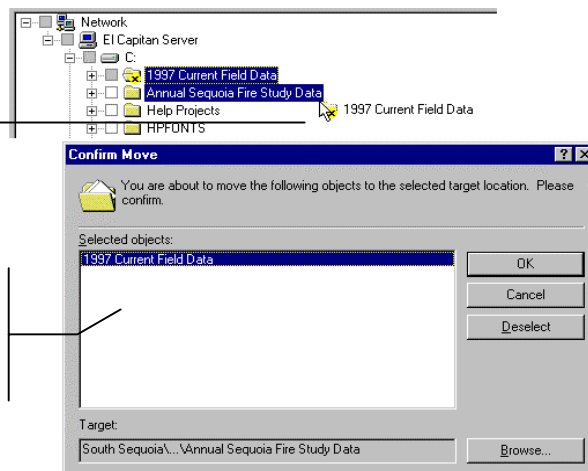
Restoring Files to a Different Folder

You can also restore files to different folders. When TapeWare restores the file, it creates a new file in the new location. Similarly, you can restore folders in new locations as well. This is useful in order to prevent overwriting files and folders that currently exist on disk.

To restore a file to a different folder, drag the file in the tree view area to the new folder. Alternatively, highlight the file and select **Move...** from the **Shortcut** menu. In the **Confirm Move** window, select a target location. TapeWare will move the file to the location you specify in the **Target** field.

To restore a folder or file to a different folder, select it and then drag it to the new folder...

...and then confirm the move in the **Confirm Move** window.



You can also restore folders and volumes in new locations. The contents of these containers move with them and are restored, along with the folder or volume, in the new location.

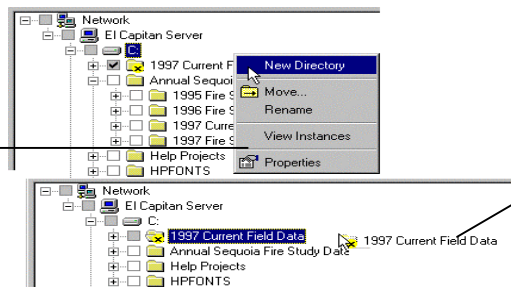
Note that when you move an instance on the **Selection** tab of a restore job, the changes you make are only recorded for that restore job. Only the current restore job will assign the file or folder the new location. When you create a new restore job, you will see the files and folders in their original locations. Likewise, the **Database** tab will continue to display files in their original locations.

Restoring Files to a New Folder

You can also create a new folder and restore files to that new folder. When TapeWare restores the files, it creates the new folder and restores the files you specified to that new location. Similarly, you can restore folders and their contents in new folders you create.

To create a new folder in which to restore the file or folder, first highlight the location you want to create the new folder in the tree view area. Then click on the **New Object** button on the **Selection** tab's **Toolbar**. Or use the **Shortcut** menu and select **New Directory**. TapeWare will create the new folder in the location you specified. Give the folder a new name and then drag into that folder the files and folders you want restored in it.

To restore a file or folder to a new folder, first highlight the container in which you want to create the new folder and then, clicking the right mouse button, select **New Directory** from the **Shortcut** menu.



Next, drag the folder or file you want to restore to the new folder you created.

Note that any new folder you create on the **Selection** tab of the restore job is only created in the restore job currently opened. Only the current job will show this new folder. When you create a new restore job or open another restore job, the new folder you created in the current job will not be displayed. Likewise, the new folder you created will not be displayed on the **Database** tab either.

Verify Selection Concepts

You select files for verifying the same way you select files for restoring, including selecting which instance to verify.

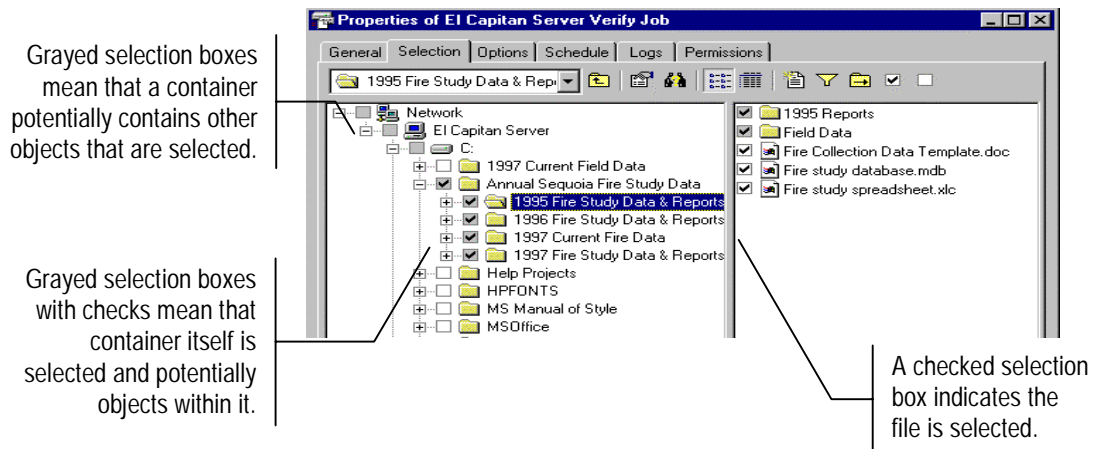
Files are selected for verifying in two steps. In the first step, the appropriate files are selected by marking them with a check and selecting the appropriate instance. In the second step, these files are filtered using multiple selection criteria.

Tip You can quickly check to see if a file was verified when it was backed up by opening the **Instance** window for that file. The **Status** field will show either **Verify complete**, **Verify failed**, or **Not verified** depending on whether or not that file was successfully verified or not when the job was run.

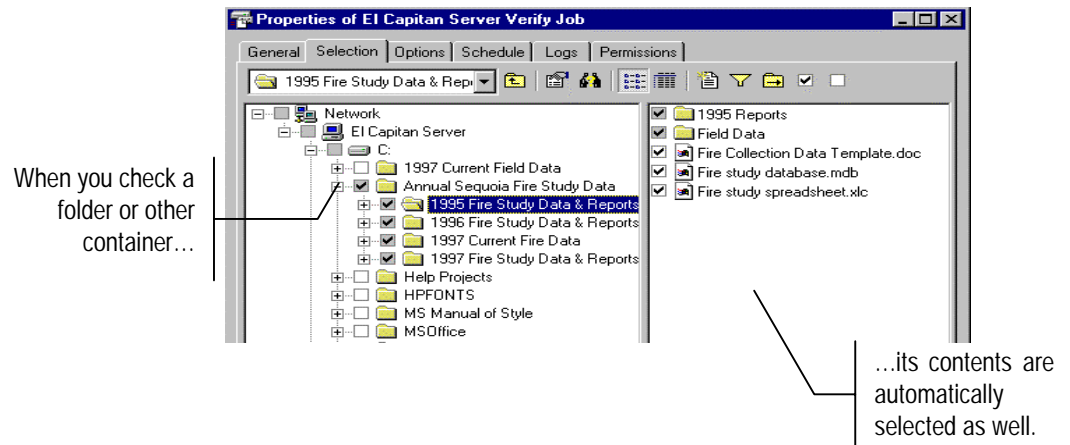
Selecting Files for Verifying

A file is selected for verifying when a check mark appears in the selection box next to the file.

When the box next to a folder or other container is gray, that means that although this folder or container is not selected itself, it potentially contains files that are selected. When the box next to a folder or other container is *grayed and checked*, the folder itself is selected and it potentially contains files within it that are selected.



You can select or unselect a file for verifying by marking or clearing the selection box next to the file. You can also mark the folder that contains the file, and not the file itself. Notice that when you mark a container, such as a folder or volume, all of its contents, including all the folders and containers it contains, are also marked.



Selecting Instances of Files

When you select a file for verifying, you can select which instance of the file you want to verify. TapeWare automatically selects the **<Latest>** instance of a file, but you can select another instance by opening the **Instances** window. The **Available instances** field shows a list of the instances of the file and the media on which those instances are stored. Select which instance you want to verify by highlighting it.

If you select the **<Latest>** parameter, TapeWare will verify the most recent instance of that file or folder.

In general, if you want to verify a specific instance of the file, you must select that file directly and specify which instance you wish to restore in the **Instances...** window.

Note that when you specify an instance date for a container, such as a folder or a volume, only those files which have matching instance dates will be selected. If a file does not have an instance date that matches the date of the container, it will not be selected. On the other hand, if you want only certain files to be verified, you can specify the instance date of a container to only select those files with matching instance dates.

Selecting Folders vs. Selecting Files

You can select the contents of the folder in one of two ways: either by individually marking the selection box of each object in that folder one by one, or by marking the selection box of the folder itself. Which method you choose is important because it effects which files TapeWare includes in the selection list *after changes have been made to that folder*.

For example, if you select a folder for verifying by marking its selection box, all of the contents of that folder are verified. If a new backup job is run before the verify job is run, TapeWare selects files for verifying using the new folder's contents. So for example, if a new file is created in that folder, TapeWare will also restore that file. Additionally, if you have selected the **<Latest>** instance of the folder, TapeWare will use the latest instance of each file in its storage management database. These files may be newer than the files you originally selected.

Selecting Files and Instances with Filters

You can also sort files for verifying with filters by clicking on the **Selection Filters** button on the toolbar and specifying which file types to include or exclude.

The selection filters you specify are applied to all of the volumes, folders, and files that have been marked for verifying. *You cannot apply different filters to different folders or volumes.*

Filters Exclude, not Include

The selection filters exclude files by filtering out files that do not meet the selection criteria. If a folder or other container has been marked for verifying, TapeWare uses the selection filters to sort through the files and to unmark any files that do not meet the selection criteria. *TapeWare does not use the selection filters to add files to the verify set.*

Selecting Files for Verify Jobs

◆ To select instances of files for verify jobs

1. Open the property sheet of the verify job and click on the **Selection** tab.
2. Mark the selection boxes next to the folders or containers you wish to include in the job.
3. In the **Instance** window that appears, select the appropriate date of the file instance you wish to verify.
4. Click on the **Selection Filters** button on the Toolbar and specify a filter selection criteria. (You can skip this step if you don't wish to apply any selection filters.)
5. Examine the tree view area and object detail area on the **Selection** tab to check that the files you intended are marked for verifying.

Selecting Specific Instances

Each time you mark a selection box of a file or folder for verifying, TapeWare automatically shows the **Instances...** window. Use this window to select which instance of the file you wish to verify.

You can also specify a particular instance of a file or folder by highlighting it in the tree view area or object detail area and then clicking the **Select Instance** button on the toolbar. TapeWare will show you the **Instances...** window with a list of the available instances for that file.

Applying Filter Criteria

When you click on the **Selection Filters** button, the **Selection Filters** window appears. This window has multiple selection filters you can use to sort through the files you have selected for verifying.

The Verify Job
Selection Filter
window

Note that each filter criterion works independent of every other. In order to be selected for verifying, each file must pass every filter criterion specified. For example, if you specify that every file selected for verifying must have been created after January 1, 1997 and must have .doc as its extension, TapeWare will only select files which meet *both* selection criteria.

The Filter Selection Criteria

This section contains a brief description of each selection filter TapeWare applies to the files and folders marked for verifying.

Note that the **Selection Filters** window for verify jobs is similar to the **Selection Filters** window for backup jobs. This allows you to use the same filters to select the files for verifying that you used for selecting files for backing up previously. This allows you to create a verify job that selects the same files as a backup job, no matter how widely distributed over the network these files may be.

Backup Range

When a file is backed up, TapeWare stores in its storage management database the date the file was backed up. Each time a file is backed up, the instance date for that file is updated to match the date of the backup. You can use this information to filter files for verify jobs.

Modify Range

Each time a file is modified, its modified date is updated. You can use this filter to verify files with a modify date that matches your criteria. TapeWare checks the directory information on the volume to see if the file should be included in the verify job. For example, you can select those files that were modified after a certain date and time or, alternatively, those that were modified before a certain date and time.

Create Range

When a file is first created, it is assigned a create date. You can use this filter to select only those files which match your criteria. TapeWare checks the created date for each file stored in the directory of the volume and uses this to select files for verifying.

Access Range

Each time a file is read, whether or not it is modified, its access date is updated. You can use this information to select files for verifying. For example, you might want to verify only those files which have been accessed (opened or read) in the past two months.

Size Range

This filter lets you select files for verifying according to their size. You might want to select only smaller files, larger files, or files between two sizes.

Instance Range

You can use this filter to instruct TapeWare to select files according to the number of instances that exist in the storage management database. You might, for example, instruct TapeWare to verify all of the files for which there is only one instance. When you set the **Instance range** filter to **At most 1**, TapeWare only verifies those files which have a single instance.

Must Match

TapeWare lets you include files using wildcard matches. If the file matches the wildcard indicated in the **Must match** field, TapeWare includes it in the verify set. For example, if you enter “*.exe” TapeWare will only verify those files with the .exe file extension.

You can specify multiple wildcards by separating each with a semicolon, “;”. For example, if you enter “*.exe;*.doc” in the **Must match** field, TapeWare selects all files that *either* have the .exe extension *or* the .doc extension.

Cannot Match

This wildcard field works just like the **Must match** field except that it *excludes* any files that match the wildcards. You can specify multiple wildcards by separating them with a semicolon; if you specify multiple wildcards, TapeWare excludes any file that matches any one of the wildcards you specify.

Wildcard Type

You can use one of three types of wildcard formats: DOS, Long, or UNIX. Select which wildcard format you wish to use from the list box.

Required Attributes

Operating systems track certain features of files called attributes that they use to manage these files. You can use these same attributes as a selection filter. In the **Required Attributes field**, if an attribute is checked, TapeWare only selects those files which have these attributes. For example, if you check **Hidden**, TapeWare only selects those files which the operating system has assigned the **Hidden** attribute.

You can select multiple attributes. In this case, only files which have all of the specified attributes will be selected.

Exclude Attributes

This field works like the **Required attributes** field except that TapeWare excludes files that match these attributes. For example, if you have checked the **Execute Only** box, TapeWare will exclude from the verify job any files with the **Execute Only** attribute.

You can select multiple attributes. A file that has any one of these attributes will be excluded. For example, if you mark the **Hidden** and **System** attributes, any file that has *either* attribute will be excluded.

Parents

When this option is checked, TapeWare verifies the directory information for any selected folder or volume. For example, if you have marked a folder, TapeWare will verify that folder only if this option is checked. When this option is not checked, directory information about folders and volumes is not verified.

Children

When this option is checked, TapeWare verifies files. When this option is unchecked, TapeWare does not verify files. This is useful if you want to verify a complex directory structure, but not the files in that directory. To verify a directory structure but not the files (children) stored in the directories, begin by marking the directory for verifying. Then clear the **Children** option. TapeWare will verify the directory structure only.

Media

TapeWare tracks instances of files and the media on which those instances are stored. You can use this information to sort files according to the media on which they appear. Only files with instances on the media in the **Media** field will be selected for the verify job. For example, if you select media named “Daily Set:1,” TapeWare will only include files in the verify job which have a valid instance on the on the media named “Daily Set:1.”

To sort files according to the media on which they appear, click the **Add...** button and select the media from the **Browse** window. Note that you must select a Media object, not a Media Folder or User/Group Folder. If there are multiple media shown in the **Media** filter field, only files which have a valid instance on *all* the media listed will be selected.

Scheduling Jobs

TapeWare offers flexible job scheduling. For ease of use and maximum security, you can use one of TapeWare's built-in job schedules. Or, alternatively, you can customize a job's schedule to efficiently meet your specific needs.

In This Chapter

- Overview
- Backup Job Scheduling Concepts
- Selecting a Backup Schedule
- Customizing Schedules
- Manual Backup Jobs
- Scheduling Restore Jobs
- Scheduling Verify Jobs

Overview

This chapter covers the **Schedule** tab of a TapeWare job, which controls when and how often a job is run.

Although the **Schedule** tab is relevant to all job types, the **Schedule** tab is especially important for backup jobs. The **Schedule** tab allows you to set up a comprehensive backup program many years into the future or, alternatively, to run a job just once or occasionally.

The first part of this chapter reviews scheduling concepts for backup jobs, while the second part explains how to create and modify backup job schedules. The final part of this chapter reviews these concepts for restore and verify jobs.


Backup Job Scheduling Concepts

This section explains some concepts and terminology helpful for choosing and selecting backup job schedules. It reviews rotating media, backup mode, and backup job type.

Rotating Media

Backup jobs performed for disaster protection are often run daily. Rather than use new media each time a job is run, TapeWare recycles or reuses the oldest media. This is efficient because it keeps costs down by limiting the amount of media needed, while still providing for data security.

The process of recycling or reusing media is referred to as **rotating media**. When TapeWare rotates media, it recycles the oldest media by overwriting it with new data. For example, suppose three backup jobs were run on Monday, Tuesday, and Wednesday. On Thursday, TapeWare will overwrite the oldest tape, in this case, Monday's tape. On Friday, Tuesday's tape, now the oldest tape, will be recycled.



Sets A rotating group of media is referred to as a **set** or **rotation set**. The media in a set are rotated, with the oldest being overwritten when a new job is run.

Set Name On backup jobs where TapeWare controls the rotation, TapeWare gives the sets **set names**, such as **Daily**, **Weekly**, **Monthly**, or **Yearly**. So, for example, when a new job is run needing a daily tape, TapeWare uses the oldest tape from that set, in this case from the **Daily** set.

Set Count Each set has a specific number of media in the set; this number is referred to as the **set count**. For some sets, the set count is as few as two, while for others it can be ten or more. For example, the **Yearly** set might consist of two media sets, while the **Daily** set might consist of ten media sets.

A group of rotating media is referred to as a **rotation set**. The **set name** is shown for each set.

General | Selection | Options | **Schedule** | Logs | Permissions

Type: Custom Start time: 11:00 Pm End of week: Friday

Type	Count
Daily: Incremental	8
Weekly: Full	8
Monthly: Full	7
Yearly: Full	2

Legend: Completed, Warning, Failed, None, Daily, Weekly, Monthly, Yearly

August 1999

Sun	Mon	Tue	Wed	Thu	Fri	Sat
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31	Yearly			

To schedule: Drag legend to calendar or right mouse click on calendar.

Each set has a **set count**, the number of rotating media in that set.

Backup Mode

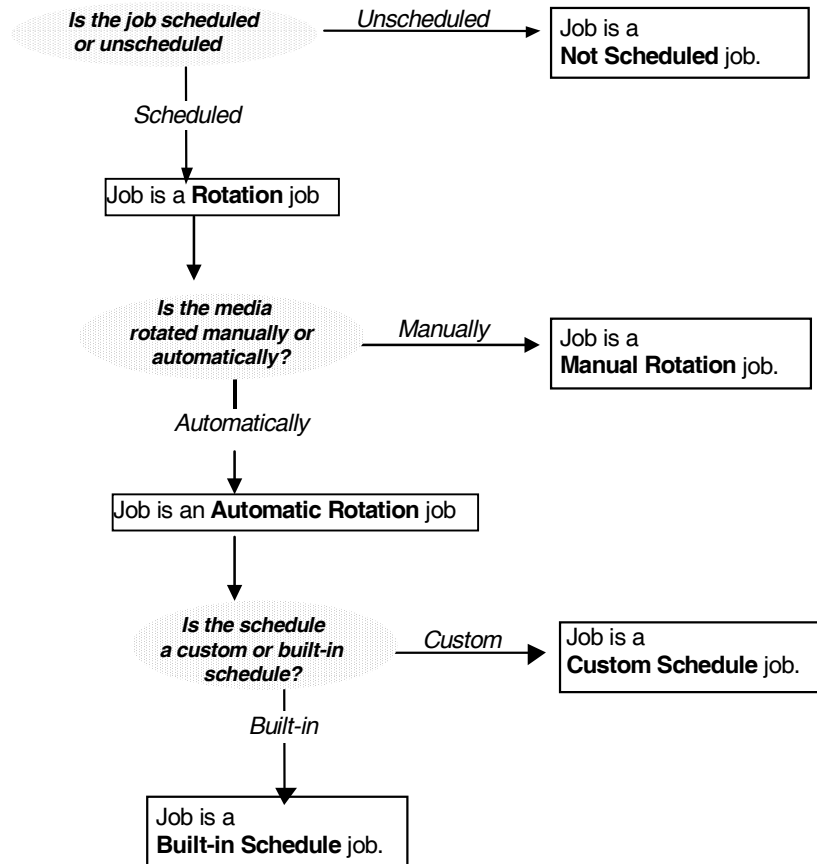
Backup jobs differ according to whether they back up all files or only changed files. Whether a job backs up all the files or only changed files is referred to as the **backup mode**.

TapeWare runs backup jobs in one of three backup modes: *full*, *differential*, or *incremental*. A job run in full backup mode will backup all files selected. A differential job backs up all files changed since the last *full* backup, while an incremental job backs up files changed since the *last* backup. Full backup jobs are the largest and take the longest to run, while incremental and differential jobs are shorter and take less time to run.

Normally, all media in a set will be created with the same backup mode. For example, all media in the **Weekly** set will be created from jobs run in full backup mode. The **Daily** set, on the other hand, is created in incremental backup mode.

Backup Job Schedule Type

TapeWare jobs can be scheduled in four distinct ways, either as *Not Scheduled*, *Manual Rotation*, *Automatic Rotation with Built-in Schedule*, or *Automatic Rotation with Custom Schedule*. The flow chart below compares the different types of schedules.



The four schedules differ from one another in the following ways:

- *Is the job scheduled or unscheduled?*

Backup jobs are either **not scheduled** or **scheduled**. Some backup jobs, such as archive jobs or historical backups, may run once or only occasionally. These jobs are not scheduled. They will run only when you instruct them to do so.

Scheduled jobs are also called **rotation jobs**. These backup jobs are generally designed for disaster protection and differ from not scheduled jobs in two ways. First, backup jobs for disaster protection are run routinely, normally daily, and thus are scheduled. Second, these jobs rotate media, while unscheduled jobs typically do not.

- *If scheduled, is the media rotated manually or automatically?*

Rotation jobs fall into two categories, **manual rotation** and **automatic rotation**. For manual rotation jobs, the user is responsible for creating the rotation sets and for selecting media for rotation (overwriting); for automatic jobs, TapeWare handles these functions by using predefined rotation sets and rotation algorithms.

Generally, most jobs designed for disaster protection use automatic rotation. Manual rotation capability increases the flexibility of TapeWare and is useful in special circumstances. Normally, however, if you wish to create and run rotation jobs, choose an automatic rotation schedule.

- *If automatic rotation, is the schedule a custom schedule or a built-in schedule?*

Automatic rotation jobs divide into two categories, **custom schedule** or **built-in schedule**.

The difference between a custom schedule job and a built-in schedule job is whether the user or TapeWare controls the backup mode and the set count. With a custom scheduled job, the user specifies the backup mode and set count; with a built-in schedule job, TapeWare handles these functions using predetermined, “built-in” schedules.

It is strongly recommended that you use a built-in schedule for disaster protection backup. These schedules are specially designed to secure your data against catastrophic loss. Data loss is possible with custom scheduled and manual rotation jobs. If a built-in schedule does not meet your needs, try making a few modifications to a built-in schedule, rather than creating an entirely new schedule.

Selecting a Backup Job Schedule

Whether or not you schedule a job and the type of schedule you select depends on several factors. Before proceeding, consider the following questions:


- What is the purpose of this job—to protect from disasters, to archive files, or to make a historical record?
- What degree of risk to the data on the network is reasonably allowable?
- Will the amount of traffic on the network require that backup jobs be scheduled to run during non-peak periods? Are there certain days of the week when running lengthy jobs will interfere with other uses of the network?

- Are there times when the tape drive will be unavailable?
- Will someone be monitoring the job as it runs?
- How large will a full backup job be?
- How much data does the media, such as a tape, hold? How much media does my budget allow me? Or, alternatively, how many tapes does my autoloader hold?

Not Scheduled Backup Jobs

Some backup jobs, such as archive jobs or historical backups, may run once or only occasionally. These jobs are not scheduled. They will run only when you instruct them to do so.

By default, backup jobs are not scheduled. To specify that a job is not scheduled, be certain that **Not Scheduled** is selected in the **Type** box on the job's **Schedule** tab.



When you want a job to run only once or occasionally, select **Not scheduled** in the **type** box on the job's **Schedule** tab.

Caution Backup jobs designed to protect data from disasters should always be scheduled. This is the most effective way of insuring that your data will be safely stored on media with regular backups.


Built-In Schedules for Backup Jobs

Backup jobs designed for disaster protection are run routinely, normally daily. TapeWare has several types of built-in schedules. Each of these built-in schedules will insure the security of your data, that is, disaster protection, the ability to recover and reconstruct data after some catastrophic loss.

The built-in schedules also provide some archive and historical backup functions, although you may design separate backup programs for these purposes. The built-in schedules vary from each other according to the number of tapes each uses and how far back historical and archival copies of files are kept. For example, some schedules use as few as four tapes, while others use as many as 30. Some schedules keep copies of some historical files available for as long as two years, while others only provide access to files from the past couple of days.

Which Built-in Schedule to Select

There are eight built-in schedules. The schedules vary in three ways: the number of days for which full data recovery is available; the level of access to historical files provided; and the minimum number of tapes or other media needed.



Full Data Recovery Period

All the built-in schedules provide for full data recovery in case of disaster. The full data recovery period is the number of days prior to the data loss for which any and every file backed up can be recovered. You will also be able to reconstruct the data for any particular day during that period.

Different schedules provide full data recovery periods for varying numbers of days prior to the last backup. For example, a GFS 30-tape job can reconstruct the data from any day of the past three weeks, while a simple 4-tape backup will only provides for reconstruction of the past two days.

Access to Historical Files

Jobs also vary according to how much access to historical files they provide. Historical files are “snapshots” of the data at particular times such as the end of the week or the end of the month.

Because they are only “snapshots,” not all files from previous weeks and months may be available. For example, you may have a historical “snapshot” of your data at the end of April 1997 and May 1997. Files that were created *and* deleted during May will not be available on either one of these historical tapes.

Different schedules will give you different levels of historical access to previous weeks, months, and years. For example, a GFS 30-tape job has 8 weekly tapes, 7 monthly tapes, and two yearly tapes. This provides the user with historical snapshots of *at least* the end of the week for the past 8 weeks, the end of the month for the past 7 months, and the end of the year for the past 2 years. On the other hand, a simple 4 tape backup will provide snapshots of only the last two end of weeks.

While these historical tapes cannot replace true historical backups, they do provide for some level of available access to historical data. Consider, for example, the yearly tapes. Each of the three GFS built-in schedules has two yearly backup tapes. The first time you run one of these jobs, you create a yearly

tape. The next yearly tape is made at the end of the current calendar year. The following year, the first tape is recycled, that is, its data is overwritten with the new data and information about the files backed up is deleted from the storage management database. This process continues, with the second tape being recycled the following year, and so forth.

Note Yearly tapes only provide you with access to files present at the time the files were backed up. No copy exists for files that were created after the oldest yearly backup and then deleted before the most recent yearly backup. To preserve a copy of these files, you must archive them before deleting them.

Minimum Number of Tapes or Media

The name of each built-in schedule indicates the minimum number of tapes or other media needed for that schedule type.

Note that the actual number of tapes needed may be more, depending on the amount of data to be backed up during a full backup and the size of your tape. If the total size of a full backup is larger than the capacity of the tape, additional tapes are required. For example, if a tape holds 1 GB, two tapes are required to back up more than 1 GB of data, three tapes to back up more than 2 GB, and so forth.

Because incremental and differential backup jobs generally have fewer selected files than full backup jobs, additional tapes may not be required for these backups. Your historical usage is the best guide to determining how many tapes these jobs will require.

If you wish to run jobs with no one monitoring the tape drive and you are not using an autoloader, be certain that you are using tapes with a capacity that is larger than the total size of the files selected for backup.

Built-In Schedules Compared

The following table compares the historical files and full data recovery capabilities of each of the built-in jobs. (The table assumes no job uses more than one tape.)

Built-In Backup Job Type	Full Data Recovery Available for Previous...	Historical “snapshots” Available for Previous...
GFS 30-tape	three 5-day weeks or two 7-day weeks (15 business days)	eight end-of-weeks seven end-of-months two end-of-years

Built-In Backup Job Type	Full Data Recovery Available for Previous...	Historical “snapshots” Available for Previous...
GFS 25-tape	two 5-day weeks (10 business days)	eight end-of-weeks seven end-of-months two end-of-years
GFS 20-tape	one 7-day week (7 business days)	six end-of-weeks six end-of-months two end-of-years
Simple 12-tape	five days	four end-of-weeks four end-of-months
Simple 11-tape	five days	four end-of-weeks three end-of-months
Simple 10-tape	five days	four end-of-weeks two end-of-months
Simple 6-tape	three days	three end-of-weeks one end-of-month
Simple 4-tape	two days	two end-of-weeks

Scheduling Backup Jobs with Built-in Schedules

To use the built-in schedules, you must select a schedule and then specify a start time, select which days of the week the job will run, and which day is the end of the week.

◆ To schedule a job with a built-in schedule

1. Select a built-in schedule by selecting it from the **Type** list on **Schedule** tab of the backup job’s property sheet.
2. Specify a time for the job to begin running in the **Start time** box. TapeWare will attempt to run the job at this time.
3. Select which days of the week you wish the job to run by clicking on the day of the week buttons at the top of the calendar on the **Schedule** tab.
4. Specify which day is the final day of the week by selecting it from the **End of week** list. On this day, TapeWare schedules **Weekly** backups.

The screenshot shows the 'Schedule' tab of the TapeWare Scheduling Jobs dialog box. The 'Type' dropdown is set to 'GFS 30-tape'. The 'Start time' is set to '11:00 Pm'. The 'End of week' is set to 'Friday'. The 'Type' and 'Count' table is as follows:

Type	Count
Daily: Incremental	12
Weekly: Full	8
Monthly: Full	8
Yearly: Full	2

The 'Legend' on the left shows status icons: Completed (green), Warning (yellow), Failed (red), and None (white). Below the legend are schedule type icons: Daily (red), Weekly (purple), Monthly (pink), and Yearly (yellow). The calendar for August 1999 shows the 30th as a green square (Completed) and the 31st as a yellow square (Warning).

Annotations:

- To schedule a job with a built-in schedule, begin by selecting the appropriate schedule
- Specify when you want the job to run in the **Start time** box.
- Specify which day you want TapeWare to run weekly backups in the **End of week** list box.
- Select which days you wish the job to run by selecting the appropriate week

Note that you can have a backup job run as frequently as everyday or only once per week. To best protect your data, select every day of the week in which new and important data is generated (that is, every business day).

Customizing Schedules

You can also create customized schedules to meet your unique backup needs. To create a customized scheduled, you can modify built-in schedules or, if needed, create an entirely new backup schedule.

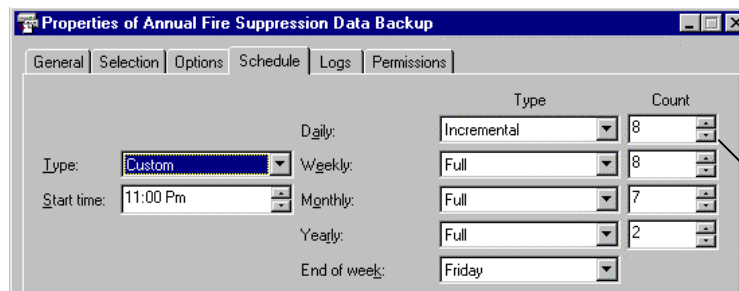
It is strongly recommended that you use the built-in schedules for disaster protection backup. These schedules are specially designed to secure your data against catastrophic loss. If you need a custom schedule, try making a few modifications to a built-in schedule, rather than creating an entirely new schedule.

Creating Custom Schedules

You can customize the schedule of a job in one of three ways: by modifying a built-in schedules, by modifying the **Custom** schedule, or by using the **Manual** schedule. This section covers how to modify built-in schedules and the Custom schedule; working with the **Manual** schedule is covered in the next section of this chapter.

Both modified built-in schedules and the Custom schedule allow you to change when monthly and weekly backups are made. However, only the **Custom**

schedule will allow you to specify the set count and the backup mode, that is, the number of tapes belonging to each type of tape (**Daily** tapes, **Weekly** tapes, and so forth) and whether those jobs are **Incremental**, **Differential**, or **Full** backups.

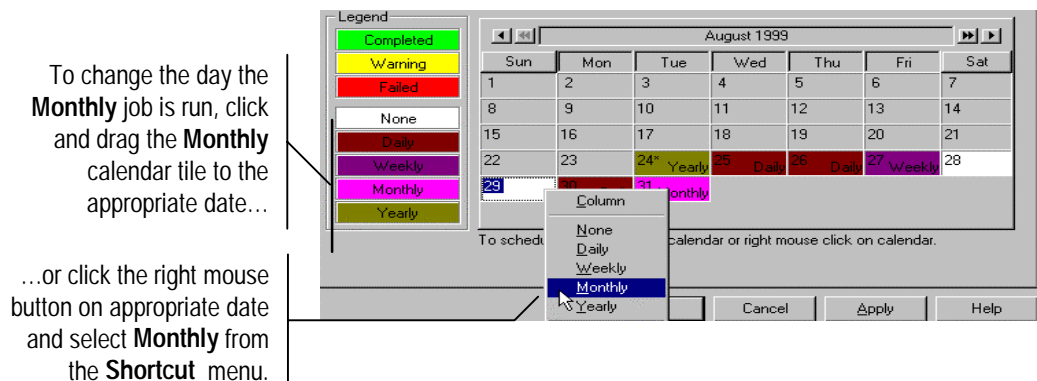


Only the **Custom** schedule lets you change the **Backup Mode Type** of each rotation set and the set **Count**.

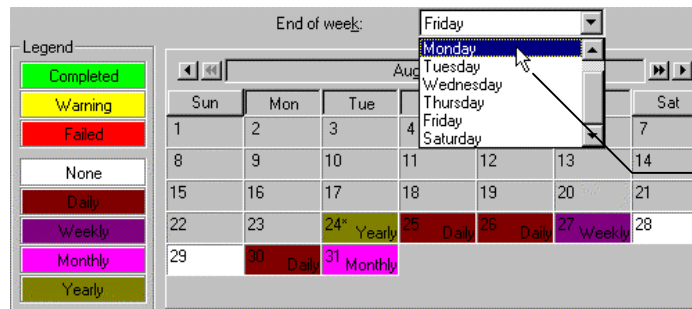
Modifying Built-in Schedules

You can change when **Daily**, **Weekly**, **Monthly**, and **Yearly** jobs are run. This is helpful when the standard schedules do not fit your particular business needs, or when a job fails to run because of a network problem or malfunction.

For example, you may wish to change the date when the **Monthly** backup is run from the last business day to the last calendar day of the month. In the following example, the **Monthly** backup was changed from the last business day of the month to the last calendar day.

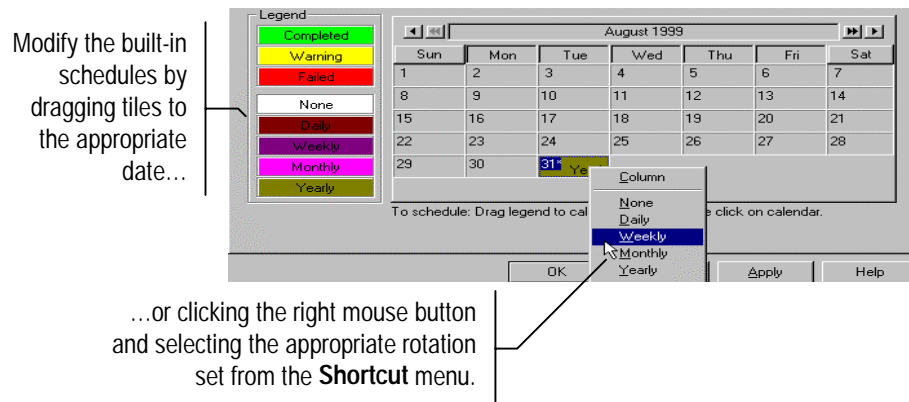


In another example, a **Weekly** backup job for some reason has failed to run when scheduled. It is important that this job be run as soon as possible or else full data recovery may not be possible. In the following example, the **Weekly** backup job that was scheduled to run on Saturday is rescheduled to run on Monday by changing Monday from a **Daily** backup to a **Weekly** backup.



You can change the day of the week weekly backup jobs are run by changing the day name in the **End of week** list box.

You modify the built-in schedules by changing the color-coded tiles on the **Schedule** tab. Drag the color coded tiles from the **Legend** to the appropriate calendar date, or click the right mouse button on the calendar and select the appropriate backup type from the shortcut menu.



Modifying the Custom Schedule

When you select **Custom** from the **Type** list on the **Schedule** tab, TapeWare sets the calendar to the **GFS 25-tape** built-in schedule. However, unlike the **GFS 25-tape** built-in schedule, the **Custom** schedule allows the user to indicate the backup mode of the **Daily**, **Weekly**, **Monthly**, and **Yearly** jobs and the number (**Count**) of each type of media or tape.

Changing the Tape Count

The **Custom** schedule is particularly useful when you wish to change either the length of the full data recovery period or the level of access to historical “snapshots.” By modifying the tape count, you can lengthen or shorten the period in which full data recovery is available or the period of time for which historical “snapshots” are available.

For example, you may wish to increase the number of yearly historical tapes available from two to three or more. To do so, change the number in the **Yearly Count** box on the **Schedule** tab to the desired number of historical backups.

	Type	Count
Daily:	Incremental	8
Weekly:	Full	8
Monthly:	Full	7
Yearly:	Full	8

With the **Custom** schedule, you can modify the **Yearly** rotation set **Count** in order to increase the number of historical tapes.

Alternatively, you may wish to expand the full data recovery period, while limiting the number of historical backups. In the following example, the full data recovery period has been expanded to four 5-day weeks (twenty business days), with only limited monthly and yearly historical tapes.

	Type	Count
Daily:	Incremental	16
Weekly:	Full	8
Monthly:	Full	7
Yearly:	Full	2

In this example, the full data recovery period has been extended to twenty business days, with limited historical jobs.

Changing the Backup Mode

TapeWare runs backup jobs in one of three backup modes: *full*, *incremental*, or *differential*. A job run in full backup mode will backup all files selected. An differential job backs up all files changed since the last *full* backup, while a incremental job backs up files changed since the *last* backup. Full backup jobs are the largest and take the longest to run, while incremental and differential jobs are shorter and take less time to run.

To change the backup mode, select the new backup mode from the appropriate **Backup Mode Type** list box on the **Schedule** tab of the job.

Determining the Length of the Full Data Recovery Period

Full reconstruction of data can be accomplished in two ways. The first method requires the most recent full backup tape and all of the incremental backup tapes since the full backup tape. The second method requires the most recent full backup tape and a differential job from the previous day.

For example, to reconstruct the data for a Wednesday, you will require one of two sets of tapes: *either*, the full backup tape from the previous end of week and all of the incremental tapes from that week (that is, Monday's, Tuesday's, and Wednesday's); *or*, the full backup tape from the previous end of week and the differential tape from Wednesday. (In some circumstances, the preceding full

backup tape will be a Monthly, not Weekly job.) As long as none of these tapes have been overwritten, full data recovery is possible.

The length of the data recovery period is determined both by the number of daily incremental or differential tapes and the number and frequency of full backup jobs (usually weekly jobs).

Incremental Jobs and Full Data Recovery

Incremental jobs are the shortest and smallest jobs to run, but they do pose some risk to full data recovery. The difference between an incremental and a differential backup is important because full data recovery is always available from a differential backup tape and a full backup tape, even when the differential tapes are being rotated (overwritten). However, when incremental tapes are overwritten or recycled, there is the possibility of data loss. Full data recovery cannot be guaranteed with incremental jobs.

Using incremental backup jobs to insure full data recovery after a disaster is not recommended, *unless you are using a built-in schedule*. However if your particular network needs to limit the amount of time available to run backup jobs, you can safely work with incremental jobs if you are careful to follow these guidelines:

- Have at least as many incremental tapes as there are days between full or differential backup jobs. For example, if you run full backup jobs every 5 days, have at least 4 incremental tapes.
- Never recycle incremental tapes between differential or full backup jobs. If you wish to run more than one incremental job in a row, be certain to not recycle any of the tapes used during this string of incremental jobs.

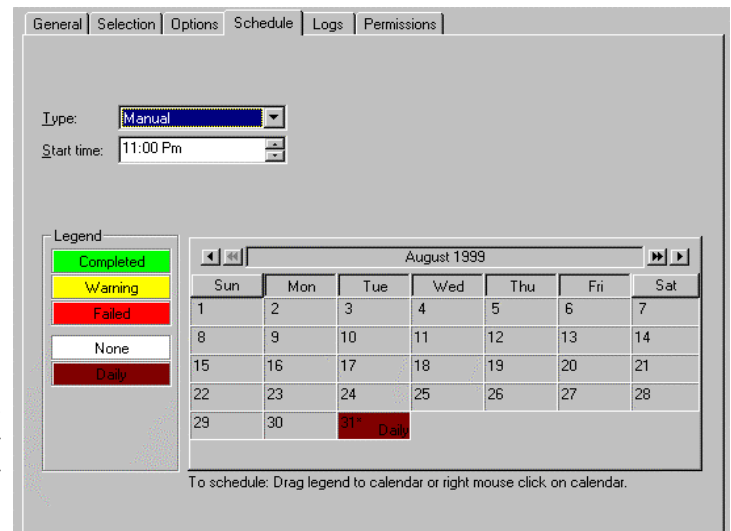
Manual Backup Jobs

Manual rotation jobs differ from unscheduled jobs because they can be scheduled. However, unlike automatic rotation jobs, the schedule does not specify the backup mode or set count.

Automatic rotation jobs allow the user to specify *in the schedule* the number of media in a set, which set is being used, and the mode of the set. For manual jobs, however, these features of the backup plan are the responsibility of the user. The user must keep track of the media sets and what media are to be overwritten. Furthermore, the user must specify the backup mode on the **Options** tab of the job.

The **Manual** schedule is provided for specialized uses; do not select this option for backup jobs intended for disaster protection.

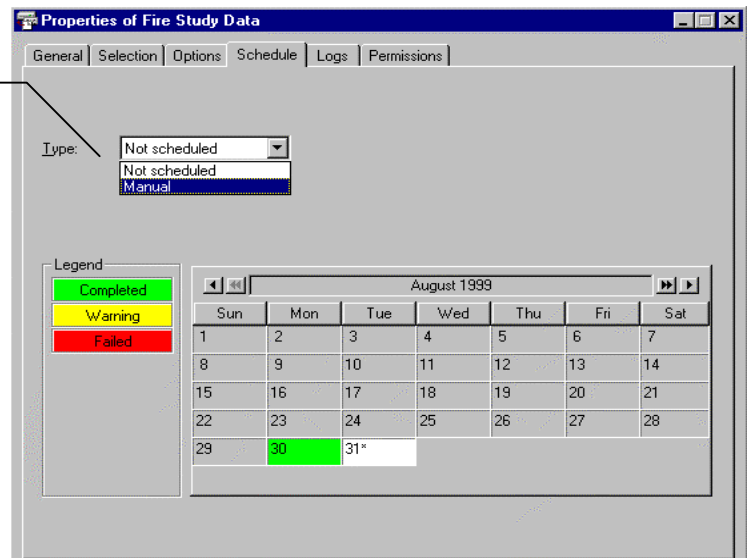
Because the user is responsible for rotating the media, the **Manual** schedule should be used only for special purposes—not disaster protection.



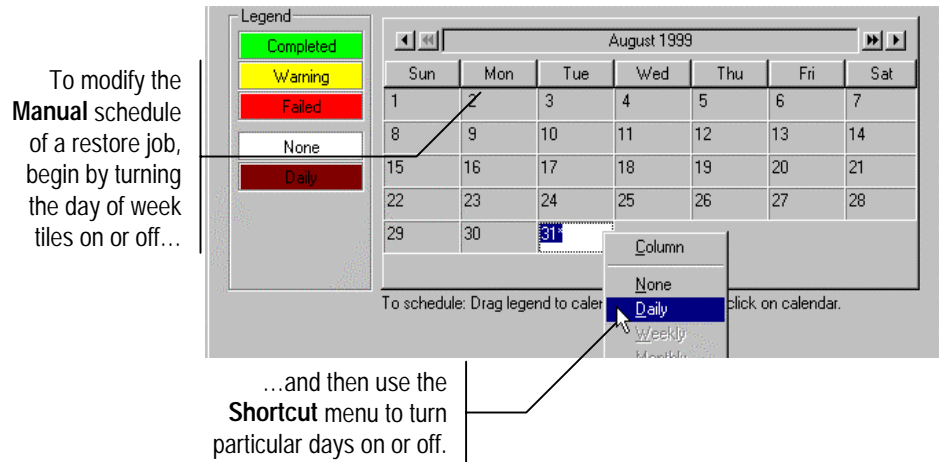
Scheduling Restore Jobs

Restore jobs are either unscheduled or scheduled with the **Manual** schedule. In general, if you want to schedule a job to run, even if it is only once, use the **Manual** schedule. If you only want your job to run when you instruct it to run, select the **Unscheduled** schedule.

There are only two schedule types for Restore jobs, **Not scheduled** and **Manual**.



If you want to run a restore job on regular occasions, select **Manual** schedule. You can specify which days it is to run by modifying the calendar. For example, you can turn off certain days of the week by clicking on the day of the week buttons on the top of the calendar. To turn a particular day off or on, click on it with the right mouse button and select either **None** or **Daily** from the **Shortcut** menu.

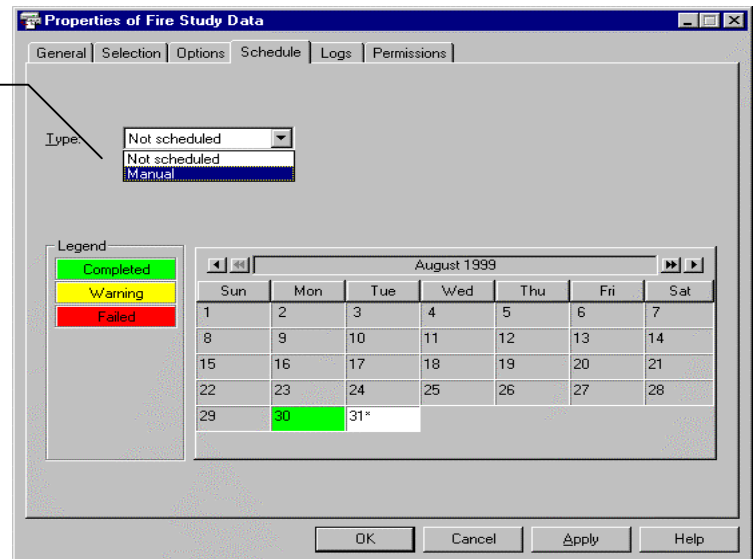


If you want a job to run only once, but need to schedule it to run during non-peak hours, try this method. Turn off all of the days of the week by clicking on the day of week name tiles. The calendar will be all white. Then click on the day you want to run the job with the right mouse button and select **Daily**. This will be the only day the job runs. Be certain to adjust the time you want the job to run in the **Start Time** box.

Scheduling Verify Jobs

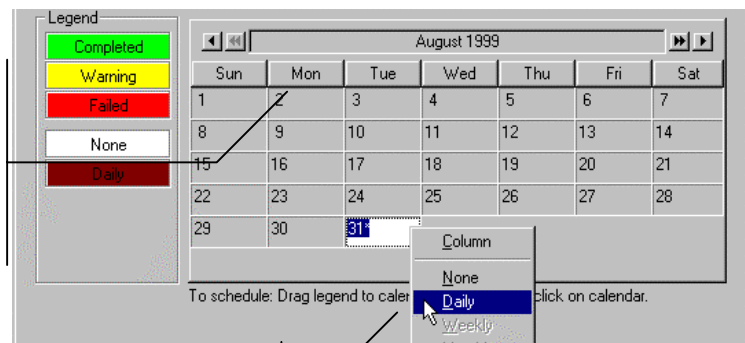
Verify jobs are either unscheduled or scheduled with the **Manual** schedule. These are the only two types of schedules available for verify jobs. In general, if you want to schedule a verify job to run, even if it is only once, use the **Manual** schedule. If you only need your job to run when you instruct it to run, select the **Unscheduled** schedule.

There are only two schedule types for **Verify** jobs, **Not scheduled** and **Manual**.



If you want to run a verify job on regular occasions, select **Manual** schedule. You can specify which days it is to run by modifying the calendar. For example, you can turn off certain days of the week by clicking on the day of the week buttons on the top of the calendar. To turn a particular day off or on, click on it with the right mouse button and select either **None** or **Daily** from the **Shortcut** menu.

To modify the **Manual** schedule of a verify job, begin by turning the day of week tiles on or off...



...and then use the **Shortcut** menu to turn particular days on or off.

If you want a job to run only once, but need to schedule it to run during non-peak hours, try this method. Turn off all of the days of the week by clicking on the day of week name tiles. The calendar will be all white. Then click on the day you want to run the job with the right mouse button and select **Daily**. This will

be the only day the job runs. Be certain to adjust the time you want the job to run in the **Start Time** box.

Job Options

TapeWare uses parameters on a job's **Options** tab to control various features necessary for running a job. TapeWare's default values are designed to be easy and secure to use, but you can modify the job option parameters to meet your particular needs.

In This Chapter

- Overview
- Backup Job Options
- Advanced Options
- Restore and Verify Job Options

Overview

This chapter covers an important tab on the property sheet of a TapeWare job: the **Options** tab. This tab controls various features of the job, such as what media is used, what backup device is used, whether files can be overwritten, and how alerts are handled.

Although this tab is relevant to all job types, the **Options** tab is especially important for backup jobs. Some types of backup schedules, such as automatic rotation jobs, treat certain job option parameters differently than do others, such as manual rotation jobs.

The first part of this chapter focuses on backup job options, while the second part explains the advanced options parameters used by backup and verify jobs. Finally, these concepts are reviewed for both restore and verify jobs.

The Backup Job Options Tab

The **Options** tab on the property sheet of backup job controls various parameters important to how TapeWare runs a backup job. Although there are numerous choices, the default values have been chosen to provide the maximum degree of

security and ease of use. Use the default values unless your particular backup needs require different settings.

The backup job
Options tab

The following section describes each option parameter and some of the possible ways each might be used.

Backup Mode

The **Backup mode** is either **Full**, **Incremental**, **Differential**, or **Snapshot**. For scheduled automatic rotation jobs, TapeWare uses the backup mode type for each backup set as indicated on the **Schedule** tab; for unscheduled or manual jobs, TapeWare uses the parameters set by the user. For more information, see “Backup Options Automatically Updated” below.

The **Backup mode**
options list box.

Full This parameter instructs TapeWare to back up all selected files.

Differential This parameter instructs TapeWare to back up all selected files that have changed since the last full backup.

Incremental This parameter instructs TapeWare to back up all selected files that have changed since the *last* backup.

Snapshot This parameter instructs TapeWare to back up *all* selected files, but it has no effect on any future scheduled job. (A snapshot backup job does not reset the archive bit after backing up all the selected files.) Use this option when you


wish to make a record of files or systems at a particular time, but do not wish to disrupt the normal backup schedule.

For more information on the differences between incremental and differential backup jobs and the importance between them for data recovery, see the previous chapter.

Write Mode

The **Write mode** is either **Append to all**, **Append to first overwrite others**, or **Overwrite all**. For scheduled automatic rotation jobs, TapeWare defaults to **Overwrite all** mode; for unscheduled and manual jobs, TapeWare uses the parameters set by the user. For more information, see “Backup Options Automatically Updated” below.

The **Write mode** options list box.



This mode determines whether the old data on the media is *overwritten* with new data, or whether the new data is *appended* to the end of the old data. When media is overwritten, all of the data previously stored on it is lost. Appending data will preserve the old data.

Unless the media is meant to be stored permanently, select **Overwrite all**. This is because when tapes or media are rotated (reused), TapeWare overwrites it. If you have appended data to the media, overwriting will result in the loss of not just the oldest material, but all of the data on the media, including the most recent. For this reason, use **Overwrite all** for media you wish to reuse through rotation, such as tapes that are part of a set of daily incremental backups, and **Append to all** or **Append to first, overwrite others** for media meant for permanent storage.

Appending is useful if the number of tapes is limited or if the tapes are several times larger than the size of the job. For example, a one gigabyte tape could hold the contents of four jobs that are less than 250 Mbytes if these jobs were appended. However, if overwrite mode is selected, only one job will be stored on a tape at once. Similar comments apply to other types of media.

Append to all This parameter instructs TapeWare to append all data to the end of the media. No data is overwritten. Select this parameter for permanent storage.

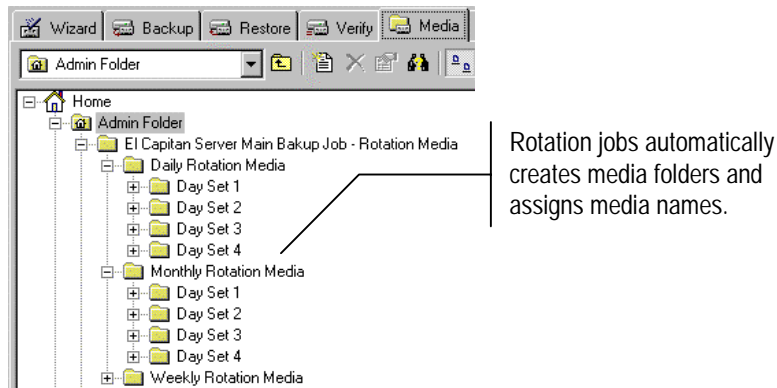
Append to first, overwrite others This parameter instructs TapeWare to append data to the end of the first media, but to overwrite all media that follows. For example, TapeWare will not overwrite the first tape inserted, but will overwrite the second, third, and later tapes. This parameter is useful if you have a set of media with old data you no longer need. By selecting this option, TapeWare

preserves your most recent data on the first media, but overwrites older, unneeded media.

Overwrite all This parameter instructs TapeWare to overwrite all media. All data on media that is overwritten is lost. Use this option for tapes that are going to be recycled.

New Media Name

The **New media name** is the name TapeWare gives to any new media it creates while running the job. For scheduled automatic rotation jobs, TapeWare automatically updates this parameter to match the media's place in the rotation schedule. For example, if the media is the first media in the yearly rotation set, TapeWare names it "Year Set 1:1". (This means the media was used with a Yearly backup job and that it was the first tape in the first set.)




For manual rotation and unscheduled jobs, TapeWare assigns any new media it creates the name listed in this field. This is also true for automatic rotation jobs that are "forced" to run. For more information, see "Backup Options Automatically Updated" below.

Auto Format Mode

The **Auto format mode** is either **No auto format**, **Auto format all**, or **Auto format blank**.

The **Auto format mode** options list



This mode determines whether or not TapeWare will format the tape automatically. Before data can be written to a tape, the tape must be formatted. When tapes are formatted, any data on them is lost. Tapes and other media are formatted when TapeWare does not recognize the media, that is, when it has no information in its storage management database about that particular media. This

will occur when the tape is blank, it has been erased, it is first used, or it has been deleted from the storage management database.

Auto formatting speeds up jobs and allows jobs to be run with no one attending them. On the other hand, disabling auto format can help insure that no data is lost by accidentally formatting a tape.

No auto format When this parameter is selected, if TapeWare encounters media that needs to be formatted (either blank or unrecognized media), it sends an alert to the alert window. While waiting for a user reply, TapeWare scans the network for devices with the media it was expecting.

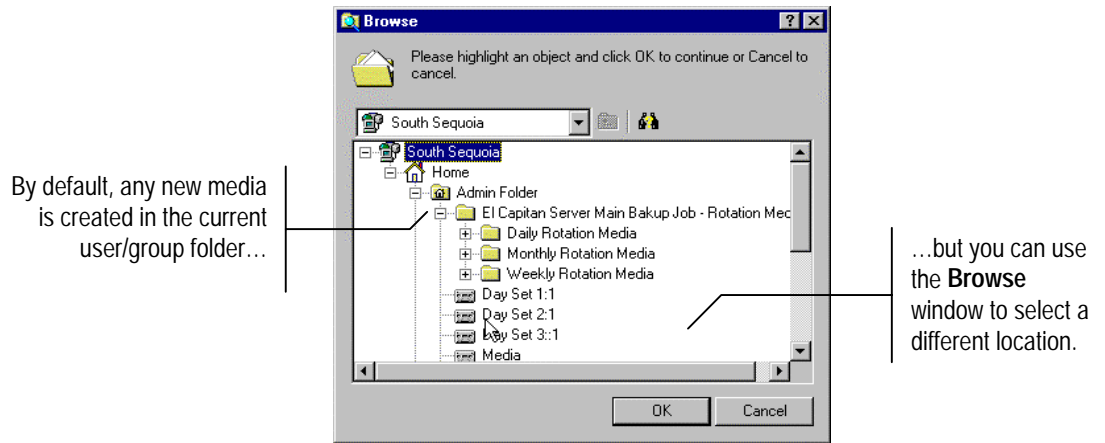
Auto format all This parameter instructs TapeWare to automatically format all of the media inserted into the tape drive which require formatting. With this parameter selected, TapeWare will automatically format all new (or blank) tapes and all unrecognized media. Select this option when you are creating a job that will run unattended.

Auto format blank This parameter instructs TapeWare to automatically format all new or blank media. However, if TapeWare encounters unrecognized media, it sends an alert to the alert window and then scans the network for the media it was expecting. This parameter can help prevent data from being accidentally destroyed by formatting, while not needlessly querying the user before formatting a blank tape. Under almost all circumstances, setting the Auto format mode to this parameter is sufficient protection against data loss and is preferable to **No auto format**.

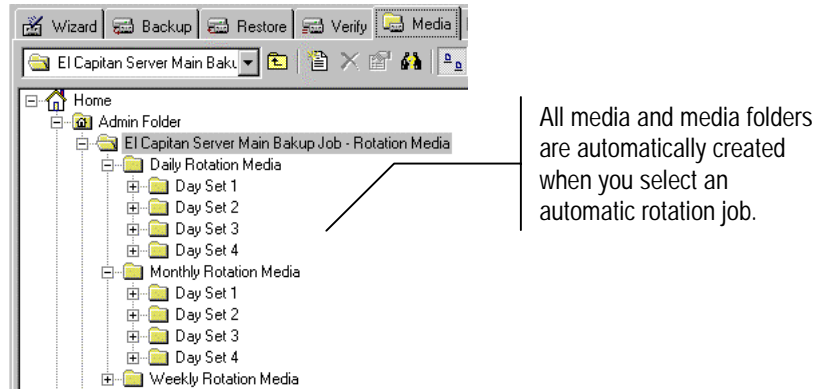
New Media Location

The **New media location** is the folder in which TapeWare will store any new tapes created while the job is running. By default, TapeWare stores media in the current User/Group folder; the media will appear there on the **Media** and **Database** tabs. You can also store the media in a Media folder inside the User/Group folder.

Select the folder in which you wish to store any new media or tapes by clicking on the **Browse...** button and then selecting the folder from the **Browse** window. If you do not already have folders set up in which to store the media, use the **Media** tab to first create additional folders.



When TapeWare runs a scheduled automatic rotation job, it automatically creates new media folders for the job. The folders are organized by the name of the job and the various rotation sets in that job. There is no reason to create these folders manually. TapeWare will automatically create these folders for you.



Number of Passes

This parameter determines how many times TapeWare will attempt to open a file in order to back up that file.

The **Number of passes** is the number of times an attempt will be made to access a file that may be in use by another user.

Number of passes:	<input type="text" value="2"/>
Delay between passes:	<input type="text" value="15"/>

The **Delay between passes** is the number of seconds between each of these attempts.

Sometimes when TapeWare is attempting to back up a file, the file may already be open, that is, currently being used by another user. Because that user may be

modifying the file, TapeWare will attempt to wait until that user is through using the file. Each time TapeWare attempts to open a file is called a **pass**. When TapeWare is unable to back up a file on the first pass, it attempts to back up that file on subsequent passes. On the last pass, TapeWare opens the file in *shared* mode.

For example, suppose the number of passes is set to 5. On the first pass, TapeWare attempts to back up all of the files in the file selection list. If, on the first pass, it encounters any files being used by other users, it skips them and attempts to back them up on the second pass. If, on the second pass, these files are still being used by other users, TapeWare schedules them for the third pass, and so forth until the final pass. On the final pass, in this case the fifth pass, any file not previously backed up is opened in shared mode, regardless if it is currently being used by other users.

The default value is **5**. If your historical usage shows that many files are open during backup or if the backup job will run when other users are likely to be working with the files to be backed up, set the value to a higher number. This will increase the number of times TapeWare attempts to open a file and may result in fewer files being opened in shared mode during the final pass.


Delay Between Passes

This determines how many seconds TapeWare waits before attempting the next pass. If your historical usage suggests that many files are being opened on the last pass, consider increasing this parameter to a higher number.

Log option

The **Log option** is either **None**, **Log only failed**, **Log only completed**, or **Log all**.

The **Log options** options list box.



TapeWare keeps a log of which files are backed up while running a job. After a job is run, you can view or print the log to see which files were successfully or unsuccessfully backed up. The default value is **Log only failed**, which writes in the log any files which were not successfully backed up. This is particularly useful for locating any problems with running the backup job.

Note: This parameter instructs TapeWare to not keep a log of the backup job as it runs.

Log only failed: This parameter instructs TapeWare to log the name of any file selected for backup but, for some reason, not backed up. Use this option to check if a backup job is running correctly.

Log only completed: This parameter instructs TapeWare to log the name of any file selected for backup and successfully backed up. You might use this option to produce a list of files backed up for archive purposes, for example.

Log all: This parameter instructs TapeWare to log the name of every file selected for backup and whether or not that file was successfully backed up. You can use this option to be certain that a backup job is running correctly as you planned.

Auto Verify Mode

The **Auto verify mode** is either **Full verify**, **No verify**, or **Quick verify**.

The **Auto verify mode** options list box.



After TapeWare backs up a file onto a tape, it can check to see that this file was backed up correctly. To verify if the file was correctly backed up, TapeWare reads the file from the tape and compares it to the original file. If any discrepancies between the two files are found, the file is considered to have failed the backup.

It is strongly recommended that the **Auto verify mode** be set to **Full verify**. Verification that data has been correctly written to the tape is an essential part of a comprehensive backup program. Further, verifying the files insures that the tape and the tape drive are working correctly. Restoring data after a disaster is no time to discover that the data was incorrectly stored to begin with.

Full verify: This parameter instructs TapeWare to compare every file on tape with the original file from the workstation or file server. This default value is strongly recommended.

Quick verify: This parameter instructs TapeWare to be certain that every file backed up onto the tape is in readable condition. It does not check to see if the data is correct, only that the data stored on the tape (incorrect or not) can be read. While selecting this option can save time, it is nonetheless not recommended.

No verify: This parameter instructs TapeWare to skip the verification step. It is not recommended.

Software Compression:

The Software Compression mode controls how TapeWare compresses or maintains the compression of files and directories.

None: This parameter instructs TapeWare to write all data to the tape in a decompressed format. If the file is stored on disk in a compressed format, the file will be decompressed before writing. This option is useful if the device supports hardware data compression and the files are to be restored to a different operating system.

Standard: This parameter instructs TapeWare to write all data to tape in the TapeWare compression format. If the file is stored on disk in a compressed format, the file will be decompressed before being re-compressed by TapeWare. This options is useful if the tape device does not support hardware data compression and the files are to be restored to a different operating system.


System: This parameter instructs TapeWare to write all data to tape in the same mode it is stored on disk. If the file is stored on disk in compressed format, TapeWare will write the data in the hosts compressed format. If the file is no compressed on disk, TapeWare will store the file on tape in a non-compressed format. This option is useful if the hardware does support data compression and the files are to be restored to the same operating system.

Both: This parameter instructs TapeWare to write all data stored on disk in compressed format, but those files that are not compressed on disk will be stored in the TapeWare compression format. This option is useful if the hardware does not support data compression and the files are to be restored to the same operating system.

Device

The **Device** option specifies which tape drive or other removable media device TapeWare will use to run the backup job.

By default, the **Device** option is set to the **Network** container, and will use any devices available on the network.



You can specify a particular device to use by adding it to the device list.

By default, TapeWare sets this parameter to the network container. When running the job, TapeWare will use whatever devices it finds on the network. If there is only one device in your storage management zone or if you only have permissions to one device, there is no reason to change this parameter.

However, if there are several devices on your network and you need to select a specific device, specify which device the job should use by selecting it from the **Device** list. (If a machine has only one device, you need not select the device, only the machine.)

For further information on using multiple streams, see "Maintaining the Flow of Data," Chapter 10, and "Storage Tab," Chapter 12..

Note that TapeWare will automatically attempt to use multiple devices whenever there are multiple backup streams. Using multiple devices can speed up a backup job tremendously. However, you may wish a job to use only one specific device or exclude a particular backup device. Use the **Device** list to specify which devices a job should use.


Note: If you are using an autoloader, the physical machine will have two (or more) drivers associated with it: the device driver (or drivers) and the autoloader driver. You may specify either driver in the **Device** list. If you specify the autoloader driver, for example, TapeWare will automatically select backup devices in the autoloader. On the other hand, if you select the device driver, the autoloader will also be automatically selected.

If you wish to use a device that is not shown in the **Device** list, click the **Add...** button and select the new device from the **Browse** window.

Media

The **Media** option specifies the database folder in which the tape or other removable media are stored in the storage management database. TapeWare will look here for media to use with this job.

The default media folder is the current user/group folder...



...but you can specify another one by adding it to the **Media** list.

The default folder is the current User/Group folder. If you wish to use media from another folder, specify which folder by selecting it from the **Media** list.

If you wish to use a folder that is not shown in the **Media** list, click the **Add...** button and select the new folder from the **Browse** window.

Media Password...

When a job creates new media, you can assign that media a password. A password prevents the media from being imported into another TapeWare storage management database and can be an important part of your overall security plan.

To assign new media your job will create a password, click the **Media Password...** button and type and confirm your password.

The **Media password** window.




Note that passwords can only be assigned when media is formatted. Additionally, media passwords are only required when importing media.

Change Mode

This option determines what action TapeWare will perform when it fails to find the media it was expecting to use for a job. When TapeWare runs a job, if the job uses specific media, TapeWare scans the network for devices with that media. If it does not find the media it expects, its response is determined by the **Change mode** parameter.

The **Change mode** options list box.



Skip device: This parameter instructs TapeWare to skip the designated device and look for other devices on the network with the proper media. TapeWare continues to scan for the correct media until it is found. As a result, the job may never run if the correct media is not found.

Force to append: This parameter instructs TapeWare to append data to whatever media it finds in the designated backup device. If it cannot find the correct media, TapeWare appends data to whatever media is available. This option will insure that the job runs, if the media contains enough room to complete the job.

Prompt: This parameter instructs TapeWare to continue to scan for the expected media and to send an alert warning that the proper media has not been found. This option will not allow a job to run with any other media except with the expected media. Additionally, this option will not search for another device that might contain the proper media.

Backup Options Automatically Updated

For further information, see "Forcing Scheduled Jobs to Run," Chapter 8.

The parameters shown on the **Options** tab fall into two categories: first, parameters that are updated automatically when TapeWare runs a scheduled *automatic rotation* job (built-in or custom), but specified manually in *unscheduled* and *manual rotation* jobs; and second, parameters that are always specified manually by the user.

When a scheduled rotation job is run, the **Backup mode**, **Write mode**, and **New media name** parameters are updated automatically.

The screenshot shows the 'Options' tab of a TapeWare configuration window. The parameters are as follows:

- Number of passes:** 2
- Delay between passes:** 15
- Log options:** Log only failed
- Backup mode:** Full
- Change mode:** Prompt
- Write mode:** Append to all tapes
- Auto verify mode:** Full verify
- Compression type:** System
- Auto format mode:** Auto format all
- New media location:** Home\Admin Folder
- New media name:** Media
- Device:** South Sequoia\Home\Admin Fo, South Sequoia\Network
- Media:** South Sequoia\Home\Admin Fo, South Sequoia\Home\Admin Fo

Buttons include 'Add...', 'Delete', 'Browse...', 'Media Password...', and 'Advanced Options...'.

For further information, see "Backup Job Scheduling Concepts," Chapter 6.

Recall from the previous chapter that TapeWare jobs can be scheduled in four distinct ways, either as *Not Scheduled*, *Manual Rotation*, *Automatic Rotation with Built-in Schedule*, or *Automatic Rotation with Custom Schedule*.

When an automatic rotation job is scheduled, the job is placed on the **Queue** tab with an indication of the date and time the job is scheduled to run. When TapeWare runs these scheduled jobs on the **Queue** tab, it automatically updates *three* of the parameters on the **Options** tab: **Backup mode**, **Write mode**, and **New media name**.

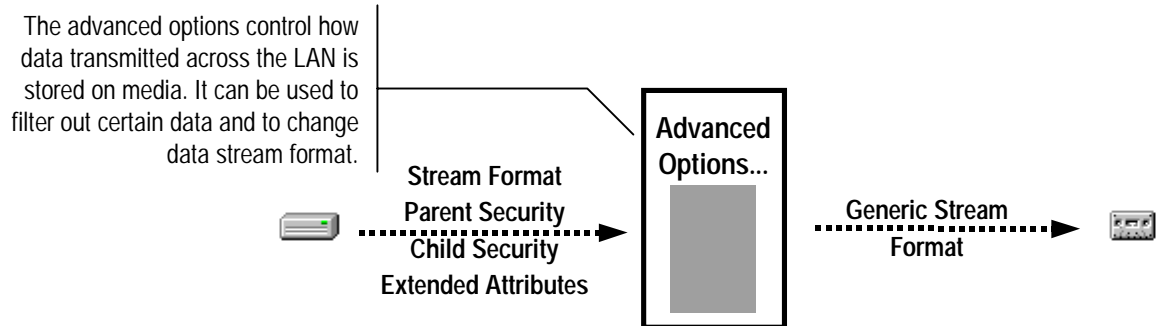
For further information on how forcing a job to run effects job parameters, see "Forcing Scheduled Jobs to Run," Chapter 8.

Note however that TapeWare does NOT automatically update these fields when you manually "force" a scheduled job to run. For example, when TapeWare automatically runs a scheduled backup job on a Monday, it changes (updates) the **Backup Mode** from **Full** to **Incremental**. But when this job is "forced" to run before its scheduled time, TapeWare does not automatically update these fields.

Note additionally that when you run an unscheduled or manual rotation job, TapeWare always uses the parameters selected by the user.


Advanced Options

The **Advanced Options...** button on the Options tab allows the user to specify certain parameters controlling how files are stored on media. In particular, these options specify whether the data on the media is stored in the same form as it was transmitted across the LAN. TapeWare can either store the data in a format specific to a particular network platform or in a generic format. Similarly, TapeWare can store all of the data it receives, or filter out some of the data used by particular network platforms or operating systems.



In general, the default values should be used. These options are provided only for advanced users who need to customize their backup jobs for unique circumstances. These options might be used in one of two circumstances: when transferring data from one network platform or OS to another; or when demands on network traffic require that a backup job be run as quickly as possible. *Unless you have specific needs that require changes to the advanced options, leave the default values unchanged.*

The **Advanced Options...** window



These options are applicable to both *backup* and *restore* jobs. Note that both job types can filter out certain data, such as security information. However, restore jobs cannot *add* data that was not originally stored on the media.

Native data streams format

Different network software transmits data across the network to TapeWare in different formats. In particular, Windows NT and NetWare use different data

stream formats. If you are going to share data from one LAN platform to another, the data should be stored on media in a common data format, *not* in the native data streams format.

Check this option when you do not plan to share data between different LAN platforms. When this option is checked, TapeWare generally runs backup jobs more quickly.

Clear this option when you plan to share data between different LAN platforms, such as from a Windows NT server to a NetWare server.

Security is an issue to consider when checking this option. When this option is checked, TapeWare backs up all security information that network software (such as Windows NT) includes in the data stream. If the option is unchecked, TapeWare uses a generic format that removes the security information.

Parent Security

When checked, TapeWare includes NetWare and Windows NT parent security information, that is, the access control list and trustee information that controls who can see and modify the *directories*. If this option is unchecked, TapeWare filters out the parent security information that it receives from the network during a backup job and that it would transmit across the network during a restore job.

Child Security

When checked, TapeWare includes NetWare and Windows NT child security information, that is, the access control list and trustee information that controls who can see and modify the *files*. If this option is unchecked, TapeWare filters out the child security information that it receives from the network during a backup job and that it would transmit across the network during a restore job.

Volume restrictions

NetWare controls the maximum amount of space a user can use on a volume. When this option is checked, TapeWare includes this information about the volume in the backup media. If this option is unchecked, TapeWare filters out the volume restrictions that it receives from the network during a backup job and that it would transmit across the network during a restore job.

Space restrictions

NetWare controls the maximum amount of space a directory can use on a volume. When this option is checked, TapeWare includes this information about the directories in the backup media. If this option is unchecked, TapeWare filters out the space restrictions that it receives from the network during a backup job and that it would transmit across the network during a restore job.

OS/2 extended attributes

When this option is checked, TapeWare includes the extended attributes for files and directories on workstations or file servers running operating systems that employ extended attributes, such as IBM's OS/2 operating system. Mac OS, Windows 95/98, and Windows NT also use extended attributes and this option also affects file servers and workstations running these operating systems. If this option is unchecked, TapeWare filters out the extended attributes that it receives from the network during a backup job and that it would transmit across the network during a restore job.

Macintosh Finder

When this option is checked, TapeWare includes the Finder information for files and directories on workstations or file servers running the Macintosh operating system. If this option is unchecked, TapeWare filters out the Finder information that it receives from the network during a backup job and that it would transmit across the network during a restore job.

NFS information

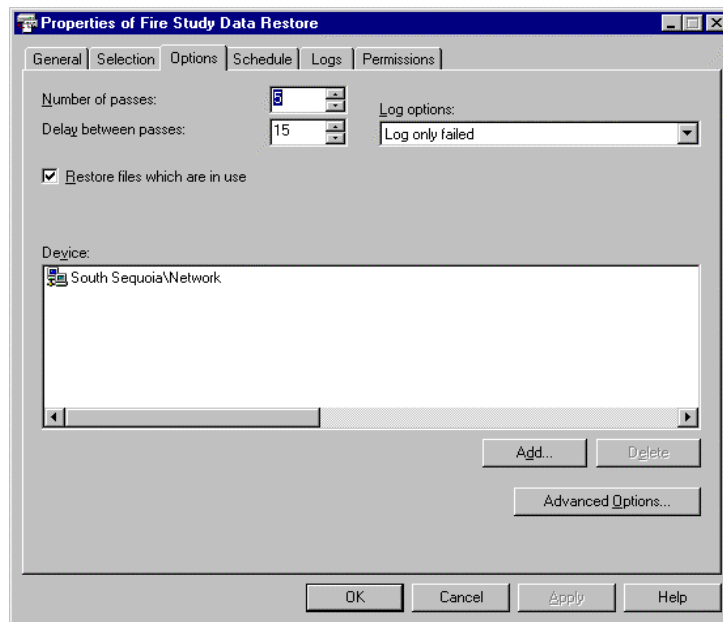
When this option is checked TapeWare includes NFS information for files and directories on workstations or file servers running a UNIX version of NFS. If this option is unchecked, TapeWare filters out the NFS information that it receives from the network during a backup job and that it would transmit across the network during a restore job.

Object owner

When this option is checked, TapeWare includes the object owner information for files and directories on workstations or file servers running NetWare. If this option is unchecked, TapeWare filters out the object owner information that it receives from the network during a backup job and that it would transmit across the network during a restore job.

Restore Job Options

There are fewer and simpler options for restore jobs. In general, these options are similar to options for backup jobs.




The restore job
Options tab.

Number of Passes

This parameter determines how many times TapeWare will attempt to access a file on the network.

The **Number of passes** is the number of times an attempt will be made to access a file that may be in use by another user.



The **Delay between passes** is the number of seconds between each of these attempts.

Sometimes when TapeWare is attempting to access a file, the file may already be open, that is, currently being used by another user. Each time TapeWare attempts to open a file is called a **pass**. When TapeWare is unable to access a file on the first pass, it attempts to back up that file on subsequent passes. For restore jobs, TapeWare attempts to open the file as many times as there are passes. When the file cannot be opened on the last pass, that file fails the restore job.)

The default value is **5**. If your historical usage shows that many files are open while a restore or verify job is being run, set the value to a higher number. This will increase the number of times TapeWare attempts to open a file.


Delay Between Passes

This determines how many seconds TapeWare waits before attempting the next pass. If your historical usage suggests that many files are being opened on the last pass, consider increasing this parameter to a higher number.

Log options

The **Log options** are either **None**, **Log only failed**, **Log only completed**, or **Log all**.

The **Log options**
options list box.



TapeWare keeps a log of which files are restored while running a restore job. After a job is run, you can view or print the log to see if the job was completed successfully. The default value is **Log only failed**, which writes in the log any files which were not successfully restored. This information can be vital for checking to see if your job ran successfully.

None: This parameter instructs TapeWare to not keep a log of the job as it runs.

Log only failed: This parameter instructs TapeWare to log the name of any file selected but, for some reason, not restored. Use this option to check if a job is running correctly.


Log only completed: This parameter instructs TapeWare to log the name of any file selected that was successfully restored.

Log all: This parameter instructs TapeWare to log the name of every file selected and whether or not that file was successfully restored. You can use this option to be certain that a job is running correctly as you planned.

Device

The **Device** option specifies which tape drive or other removable media device TapeWare will use to run the job.

By default, the **Device** option is set to the **Network** container, and will use any devices available on the network.



...but you can select a particular device by adding it to the **Device**

By default, TapeWare sets this parameter to the network container. When running the job, TapeWare will use whatever device it finds on the network. If there is only one device in your storage management zone or if you only have permissions to one device, there is no reason to change this parameter.

However, if there are several devices on your network and you need to select a specific device, specify which device the job should use by selecting it from the **Device** list. (If a machine has only one device, you need not select the device, only the machine.)

If you wish to use a device that is not shown in the **Device** list, click the **Add...** button and select the new device from the **Browse** window.

Advanced Options

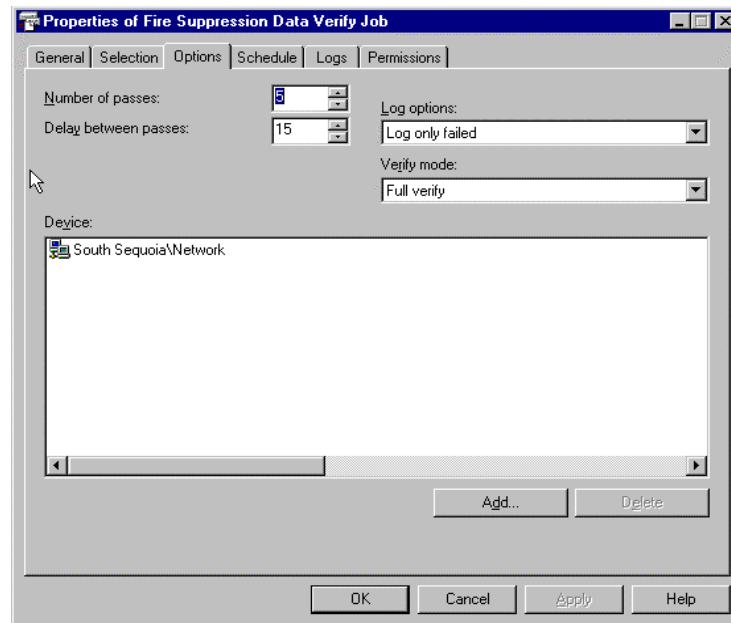
For restore jobs, you may also specify advanced options. Generally, these options work just like they do for backup jobs, that is, as filters that exclude certain types of data.

However there are some changes to the advanced options of a restore job from those of backup jobs.

- The **Native data streams format** does not function with restore jobs. TapeWare does not change the data stream format from its stored format. This option is important for backup jobs when you wish to share data from one network platform to another.
- Note that the data filters, such as security information and directory attributes, cannot add data that was not originally stored on the media. For example, if when the original backup job was run, the **Volume restrictions** box was unchecked, checking the box when running the restore job will have no effect. Because the volume restrictions are not stored on the media, they cannot be restored.

Verify Job Options

There are fewer and simpler options for verify jobs. In general, these options are similar to options for backup jobs.



The verify job
Options tab.

Number of Passes

This parameter determines how many times TapeWare will attempt to access a file on the network.

The **Number of passes** is the number of times an attempt will be made to access a file that may be in use by another user.

A close-up of the two spinners from the dialog box. The "Number of passes:" spinner is set to 2, and the "Delay between passes:" spinner is set to 15.

The **Delay between passes** is the number of seconds between each of these attempts.

Sometimes when TapeWare is attempting to access a file, the file may already be open, that is, currently being used by another user. Each time TapeWare attempts to open a file is called a **pass**. When TapeWare is unable to access a file on the first pass, it attempts to back up that file on subsequent passes.

The default value is **5**. If your historical usage shows that many files are open while a verify job is being run, set the value to a higher number. This will

increase the number of times TapeWare attempts to open a file and may result in fewer files failing to be verified.


Delay Between Passes

This determines how many seconds TapeWare waits before attempting the next pass. If your historical usage suggests that many files are being opened on the last pass, consider increasing this parameter to a higher number.

Log options

The **Log options** are either **None**, **Log only failed**, **Log only completed**, or **Log all**.

The **Log options**
options list box.



TapeWare keeps a log of which files are verified while running a verify job. After a job is run, you can view or print the log to see if the job was completed successfully. The default value is **Log only failed**, which writes in the log any files which were not successfully verified. This information can be vital for checking to see if your job ran successfully.

None: This parameter instructs TapeWare to not keep a log of the job as it runs.

Log only failed: This parameter instructs TapeWare to log the name of any file selected but, for some reason, not verified. Use this option to check if a job is running correctly.


Log only completed: This parameter instructs TapeWare to log the name of any file selected that was successfully verified.

Log all: This parameter instructs TapeWare to log the name of every file selected and whether or not that file was successfully verified. You can use this option to be certain that a job is running correctly as you planned.

Verify Mode

The **Verify mode** is either **Full verify**, **No verify**, or **Quick verify**.

The **Verify mode**
list box.



When TapeWare runs a verify job, it checks to see if the data on the media is readable and whether or not it matches data from the original source (that is, from workstations or file servers).


Full verify: This parameter instructs TapeWare to compare every file on the media with the original file from the workstation or file server. This default value is strongly recommended.

Quick verify: This parameter instructs TapeWare to be certain that every file backed up onto the media is in readable condition. It does not check to see if the data is correct, only that the data stored on the media (incorrect or not) can be read. While selecting this option can save time, it is nonetheless not recommended.

Device

The **Device** option specifies which tape drive or other removable media device TapeWare will use to run the job.

By default, the **Device** option is set to the **Network** container, and will use any devices available on the network.



...but you can select a particular device by adding it to the **Device** window.

By default, TapeWare sets this parameter to the network container. When running the job, TapeWare will use whatever device it finds on the network. If there is only one device in your storage management zone or if you only have permissions to one device, there is no reason to change this parameter.

However, if there are several devices on your network and you need to select a specific device, specify which device the job should use by selecting it from the **Device** list. (If a machine has only one device, you need not select the device, only the machine.)

If you wish to use a device that is not shown in the **Device** list, click the **Add...** button and select the new device from the **Browse** window.

Running Jobs

TapeWare will automatically run jobs that are scheduled. You can view jobs scheduled to be run on the **Queue** tab and you can track the progress of a job as it runs from the job's **Status** window.

In This Chapter

- | | |
|----------------------------|------------------------------------|
| • Overview | • The Job Status Window |
| • The Queue Tab | • Viewing and Printing the Job Log |
| • Running Scheduled Jobs | • Viewing and Printing Audit Logs |
| • Running Unscheduled Jobs | |

Overview

For jobs that are not scheduled, TapeWare will only run the job when you instruct it to do so. Scheduled jobs run automatically as scheduled. You can view what jobs are scheduled to run on the **Queue** tab, which indicates when a job is scheduled to run and provides a short summary of a job's progress as it runs. When TapeWare automatically runs a scheduled job from the **Queue** tab, it updates the option parameters for that job before running it.

You can also “force” scheduled jobs to run before they are scheduled. When a scheduled job is forced to run, TapeWare does not automatically update parameters on the job's **Option** tab. Forcing a job to run can also effect the permissions TapeWare uses when it runs the job. Be certain to read “Forcing Scheduled Jobs to Run” below for more information how forcing a scheduled job to run effects the parameters TapeWare uses to run the job.

The **Status** window provides detailed information about the progress and status of jobs as they run. You can use this window to see that a job is running properly. After a job has completed running, you can view and print the **Job Log**

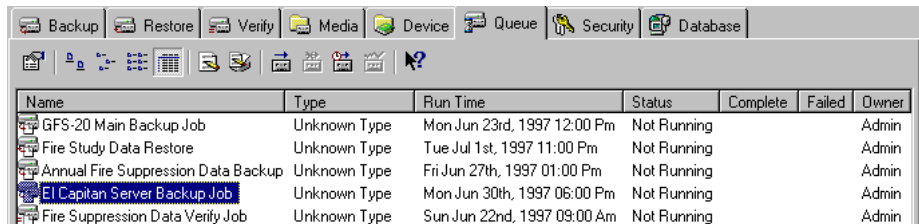
to check what files were successfully or unsuccessfully backed up, verified, or restored.

The Queue Tab

After a job has been scheduled to run, TapeWare displays the job and information about it on the **Queue** tab. This tab shows all of the jobs that are scheduled to run. As new jobs are created and scheduled, they are listed on the **Queue** tab. Each scheduled job is listed only once using information for the next scheduled time that jobs is to run.

Unscheduled jobs that TapeWare has been instructed to run also appear on the **Queue** tab, but only after you have manually instructed them to run.

The Detail view of
the Queue tab.



Name	Type	Run Time	Status	Complete	Failed	Owner
GFS-20 Main Backup Job	Unknown Type	Mon Jun 23rd, 1997 12:00 Pm	Not Running			Admin
Fire Study Data Restore	Unknown Type	Tue Jul 1st, 1997 11:00 Pm	Not Running			Admin
Annual Fire Suppression Data Backup	Unknown Type	Fri Jun 27th, 1997 01:00 Pm	Not Running			Admin
El Capitan Server Backup Job	Unknown Type	Mon Jun 30th, 1997 06:00 Pm	Not Running			Admin
Fire Suppression Data Verify Job	Unknown Type	Sun Jun 22nd, 1997 09:00 Am	Not Running			Admin

The Queue Tab Detail View

The **Detail** view on the **Queue** tab provides the most useful and important information about scheduled and currently running jobs.

The **Name** and **Type** fields show the name of the job and whether it is a backup job, restore job, or verify job. The **Run Time** field indicates the date and time the job is next scheduled to run. The **Schedule Type** field indicates what type of schedule that job has, such as **GFS-30** or **Custom**.

The **Status**, **Complete**, and **Failed** fields provide a short summary of a job's current status. The **Status** field indicates if the job is currently running and if so, what procedure is currently being performed. The **Complete** and **Failed** fields indicate what percentages of the files selected have been successfully or unsuccessfully backed up, restored, or verified.

The **Owner** field indicates what user permissions information TapeWare will use to calculate the permissions for the job when it is run.

Running Scheduled Jobs

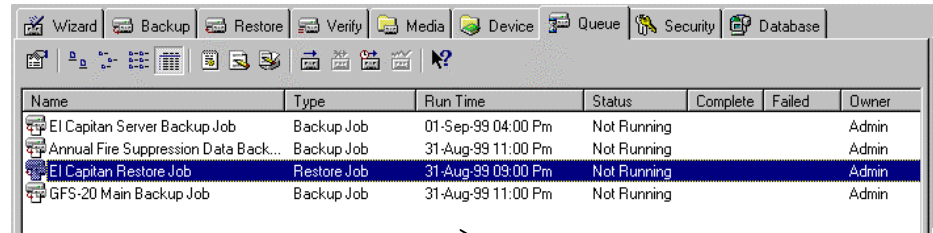
Scheduled jobs are normally run automatically by TapeWare, but you can also “force” a scheduled job to run.



The Details
button

Automatically Running Scheduled Jobs

When you close the property sheet of a job, TapeWare calculates the next time the job is scheduled to run and places the job on the **Queue** tab. The **Run Time** field in the **Details** view on the **Queue** tab shows the date and time the job is scheduled to run. This is true for jobs scheduled with automatic or manual rotation schedules.



The screenshot shows the TapeWare interface with the 'Queue' tab selected. A table lists several scheduled jobs. The 'Run Time' column shows the date and time each job is scheduled to run. The job 'EI Capitan Restore Job' is highlighted in blue.


Name	Type	Run Time	Status	Complete	Failed	Owner
EI Capitan Server Backup Job	Backup Job	01-Sep-99 04:00 Pm	Not Running			Admin
Annual Fire Suppression Data Back...	Backup Job	31-Aug-99 11:00 Pm	Not Running			Admin
EI Capitan Restore Job	Restore Job	31-Aug-99 09:00 Pm	Not Running			Admin
GFS-20 Main Backup Job	Backup Job	31-Aug-99 11:00 Pm	Not Running			Admin

The **Run** time of a job is shown on the **Queue** tab.

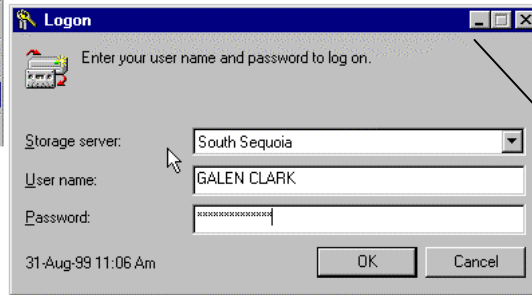
These jobs will run automatically if TapeWare is open at the scheduled date and time. Note that it is not necessary for a user to be logged on to TapeWare for the job to run. TapeWare will run scheduled jobs even if the user who created the job has logged out.

For example, suppose you have scheduled a job to run at 11:00 p.m. tonight. When you leave your workstation, log out of TapeWare. *Do not exit or close TapeWare.* When the **TapeWare Logon** window appears, click on the **minimize** button to close the window. Although no user will be logged on, TapeWare will still be open and will execute the job at the scheduled time.

Tip You can install TapeWare as a service on machines running Windows 95/98 and Windows NT. When installed as a service, TapeWare will start automatically whenever the systems starts up and run in the background without any user interface. If you want to be certain that scheduled jobs always run, consider installing TapeWare as a service. For more information, see “Installing TapeWare as a Service,” Chapter 2.



If you want scheduled jobs to run, don't exit or close, instead, log out. Select **Log out** from the **File** menu...



...and then minimize the **Logon** window, to insure no unauthorized access to the network. Another alternative is to install as a service.

Note If the storage management server is turned off when a job is scheduled, the job will run when that machine is started again. Jobs scheduled to run will begin running five minutes after you start TapeWare. (If TapeWare is run as a service, this will be five minutes after startup.) This five minute lag allows you to modify, update, or cancel any pending jobs before they run.

Security and Scheduled Jobs

Scheduled jobs will run whether you log out or not, as long as TapeWare is open. If you have not logged out, however, unauthorized users will be able work with your security clearance. For this reason, when you leave TapeWare open to run a scheduled job, be certain to log out before leaving your workstation. This is the only way to insure that no unauthorized users gain access to sensitive data.

Warning Do not leave the main TapeWare window open when you are not at your workstation. Doing so allows users without security clearance unauthorized access to the LAN. Be certain to log out of TapeWare before leaving your workstation. If you have jobs scheduled to run, log out of TapeWare instead of closing or exiting TapeWare.


Forcing Scheduled Jobs to Run



Run button.

You can “force” scheduled jobs to run prior to their scheduled time by selecting the job and clicking the **Run** button on the **Toolbar**. Alternatively, you can select **Run** from the **Job** menu, the **Queue** menu, or the **Shortcut** menu. TapeWare will execute the job immediately. Note that forcing a job to run does not affect the next scheduled run date and time.

You can force a scheduled job to run by selecting it and clicking the **Run** button...



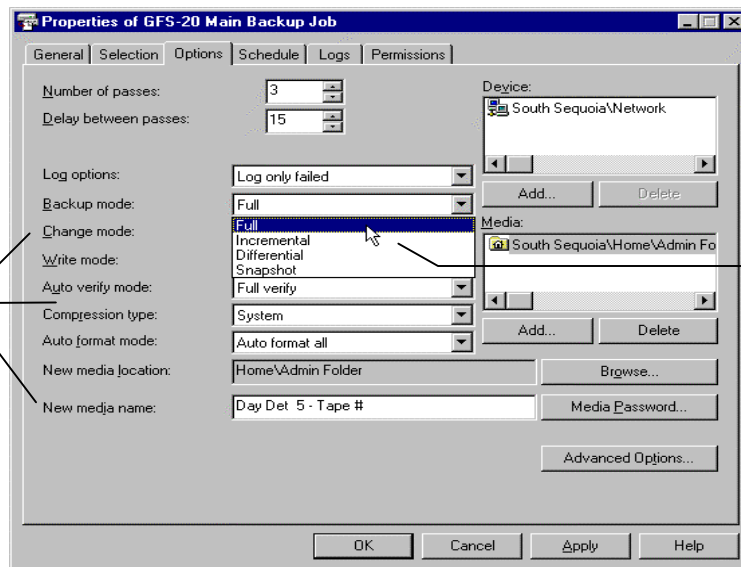
...or clicking the job with the right mouse button and selecting **Run**.

How Forcing Jobs to Run Effects Job Parameters

When you force a scheduled job to run before its scheduled time, TapeWare does not automatically update certain parameters on the **Options** tab of the job. Recall that when a scheduled job *with an automatic rotation* is run, TapeWare updates the **Backup mode**, **Write mode**, **Media folder**, and **New media name** parameters on the **Options** tab of the job to reflect that job's place in the rotation schedule. However, when a job is forced to run before its scheduled time, TapeWare does not update these option parameters.

For example, suppose that a backup job is scheduled to run as an incremental job in the evening. If it is forced to run before its scheduled time, TapeWare will not update the **Backup Mode** parameter. In this case, if the last time the job was run, it was as a full backup job, the **Backup Mode** parameter on the job's **Option** tab will still be set to **Full**. As a result, when you force the job to run, it will be run as a full backup job, even though it is next scheduled to run as an incremental job.

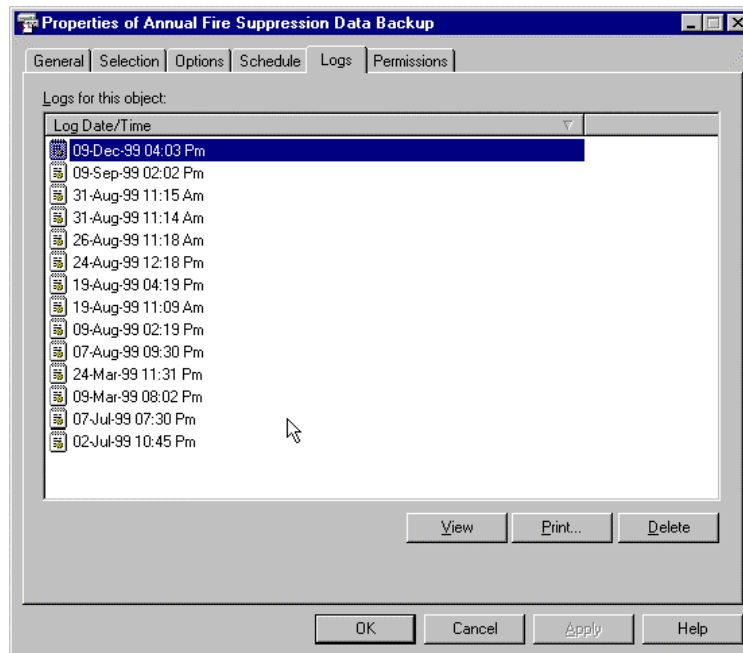
When you force a scheduled job to run, the **Backup mode**, **Write mode**, or **New media name** are not updated.



You must manually set the **Backup mode** and other options before forcing a scheduled job to run.

Forcing a job to run can be useful when a job failed to run for some reason. For example, suppose a full backup job is scheduled for a Saturday, but a LAN equipment malfunction prevented the job from being run as scheduled. It is important that another *full* backup job be run before the next *incremental* job. This is the only way to insure that the full data recovery period is not compromised. On Monday, the full backup job that failed to run properly can be re-run by forcing the job to run. Before you run the job, open the job's property sheet and check to see that the proper job type and media are selected.

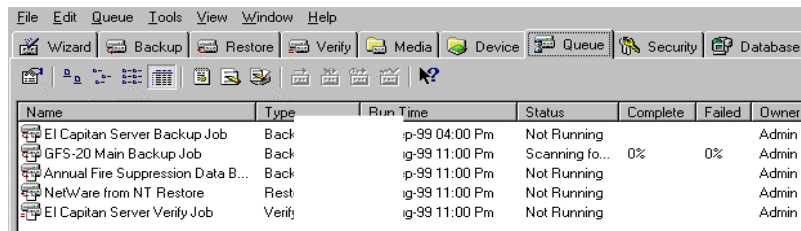
Before forcing a scheduled job to run, you should always check the **Options** tab of the job to see that the option parameters are set correctly. If you are forcing the job to run because an earlier job failed to run properly, you can look at the log of the failed job to see what parameters the job would have used.



The Logs Tab

How Forcing Jobs to Run Effects Permissions

When a job is run, TapeWare will check for the appropriate permissions to the device, files, media, and so forth. TapeWare calculates these permissions by using the permissions of the job's **Owner**. The owner of a job is the user who either scheduled the job or forced the job to be run. After a job is run, the job owner is reset to the last user who changed the job properties; forcing a job to run does not permanently change the job owner.



Name	Type	Run Time	Status	Complete	Failed	Owner
El Capitan Server Backup Job	Back	p-99 04:00 Pm	Not Running			Admin
GFS-20 Main Backup Job	Back	ig-99 11:00 Pm	Scanning fo...	0%	0%	Admin
Annual Fire Suppression Data B...	Back	p-99 11:00 Pm	Not Running			Admin
NetWare from NT Restore	Rest	ig-99 11:00 Pm	Not Running			Admin
El Capitan Server Verify Job	Verif	ig-99 11:00 Pm	Not Running			Admin

When a scheduled job is run, the job's **Owner** is used to calculate what permissions are necessary for the job to run. If you force a job to run, the job's **Owner** becomes whoever forced the job to run.

When the TapeWare administrator creates and schedules a job, the owner of the job is the **Admin**. TapeWare will use the TapeWare administrator's permissions when running the job. Similarly, if another user creates and schedules a job, that user will be the job's owner and TapeWare will calculate the job's permissions using that user's permissions.

However, if a scheduled job is forced to run, the person who forces the job to run becomes the job's new owner. So, for example, if the TapeWare administrator forces a job to run that another user has created, the TapeWare administrator becomes the job's new temporary owner and TapeWare calculates the permissions using the TapeWare administrator's permissions.

Changing the job's owner can be useful for managing security. A user can create and schedule a job, even though that user lacks the proper permissions to run that job. Another user, such as the TapeWare administrator, can then force that job to run with their own permissions.

Viewing and Printing Scheduled Job Instructions

Whenever a job is scheduled and placed on the **Queue** tab, TapeWare creates an set of *instructions* for that job. Included in a job's instructions is information about which media set must be available to be used and which backup devices it may be inserted into. For example, when running an automatic rotation job, the instructions for that job include the name of the media TapeWare is expecting to use when that job runs next, such as "Daily Set 1" or "Yearly Set 2". The instructions also include the name of the backup devices which TapeWare expects to be available when the job runs.

Tip You can use the instructions in order to insure that all of your jobs run correctly by planning ahead to see that each job has the media it requires before it runs. For example, you can print the instructions and then assign a co-worker the task of inserting the proper media into various backup devices by the required time.



The View
Instructions button

You can view the **Instructions** for the jobs currently scheduled on the **Queue** tab by clicking the **View Instructions** button on the **Toolbar** of the **Queue** tab or by selecting **Instructions** the **Tools** menu. TapeWare will open the **Instructions** in a window of an external text editor, such as Notepad. You can also print the **Instructions** by clicking on the **Print Instructions** button or from the **Tools** menu. (To change the external text editor with which you wish to view and print the **Instructions**, select **Preferences** from the **File** menu.)

You can print or view the instructions for all of the jobs scheduled on the **Queue** tab. Use the instructions to insure that the proper media and backup devices are available before a job is scheduled to run.

```

Instructions.log - Notepad
File Edit Search Help

-----
Job and Media Requirements
Generated for : South Sequoia
Generated by  : Admin
Generated at  : 31-Aug-99 11:22 AM
-----

Job El Capitan Server Backup Job by Admin
Insert rotation : Weekly Set 1
Into device    : El Capitan Server\...\Dev-0.0.6.0: Seagate STT800
                or : El Capitan Server\...\Dev-0.0.6.0: Seagate STT800
No later than  : 10-Sep-99 04:00 PM
  
```

Running Unscheduled Jobs

If you did not schedule the job, then you must manually instruct TapeWare to run the job when you want it to run. To run the job, select the job and then click the **Run** button on the **Toolbar**, select **Run** from the **Job** menu, the **Queue** menu, or the **Shortcut** menu, or double click on the job. TapeWare will execute the job immediately.

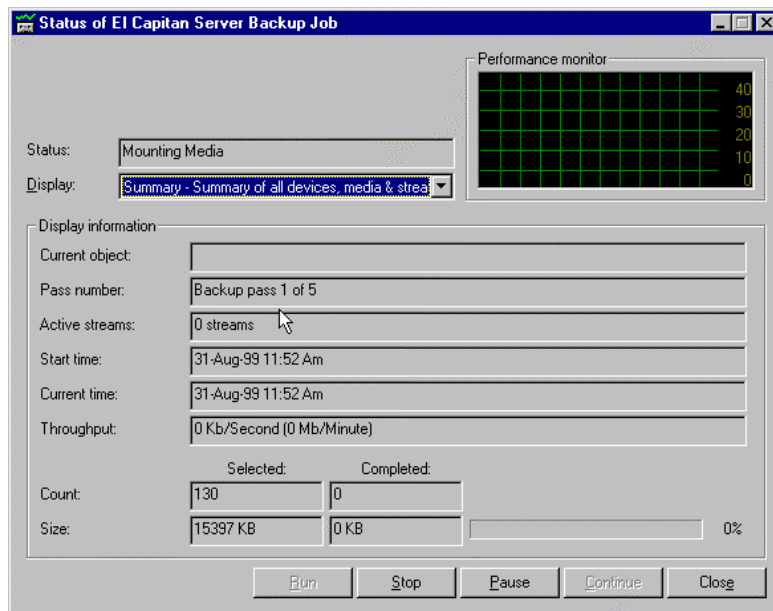
Unscheduled Job Parameters and Permissions

When an unscheduled (or manual rotation) job is run, TapeWare uses the current parameters on the job's **Option** tab.

Similarly, the owner of the job is the person that instructed TapeWare to run the job. TapeWare calculates the permissions of the job using this user's permissions, that is, the permissions of the job's owner. Note that the creator of a job and its owner are not necessarily the same user.


The Job Status Window

Each time TapeWare runs a job, it goes through a predetermined series of steps. Many of these steps are indicated in the **Status** window of the job. If you have run the job manually, TapeWare will automatically show the **Status** window. If the **Status** window of a job that is running is not displayed, you can open it by selecting the job and clicking on the **Status** button, or by selecting **Status** from the **Job** menu, the **Queue** menu, or the **Shortcut** menu.



The summary view of the **Job Status** window.

You can change what information is displayed in the **Status** window by selecting an option from the **Display** list box. This allows the user to inspect how various components of the backup job are running and locate potential problems. Many of the fields in the display box are self-explanatory; if you require additional help, use the context sensitive help. To see a short description of the different fields, click on the field with the right mouse button to view the **What's this?** window. Or, alternatively, press **F1** to get context sensitive help for the targeted area.




The **Display** list box shows the different **Job Status** window views available.

Job Status Messages

As jobs are run, TapeWare displays messages in the **Status** message box indicating the progress of the job. These messages are also displayed in the job's **Status** field on the **Queue** tab. Which messages are displayed depends on the type of job being run. The following short descriptions indicate what procedures TapeWare is performing as each messages is displayed.

Building...Selection List The first step is to create a list of files to be backed up, verified, or restored. For backup jobs, TapeWare uses the selection criteria and the backup job type (whether full, incremental, or differential) to create a list of files to be backed up. The number of files and the total size of the selected files are indicated in the **Count** and **Size** fields under **Selected**.

The **Count** and **Size** fields.



Mounting Media If tape or other media present in the device are not recognized, such as when a new tape is inserted, TapeWare displays this message as it mounts the media. During this step, TapeWare reads identification information stored on the tape. TapeWare then checks to see if the media already exists in the storage management database and whether or not the current job can be run using this media.

If after mounting the media, TapeWare recognizes it as media it can use with the current job, TapeWare proceeds to the next step.

If after mounting the media, TapeWare does not recognize it as media it can use for this current job, the next step is determined by the option specified in the **Change Mode** list box on the **Options** tab of the job.

Scanning for Device: This message is displayed when TapeWare is looking for a device to use with the current job. This message might be displayed when the current media cannot be used with this job or when TapeWare cannot find a device on the LAN.

Note Many times this message will be accompanied by an alert. You can view any current alerts by clicking the **Alert** button on the **Status** bar.

Formatting Media: If the media is unformatted, TapeWare formats the media before proceeding and displays this message. If the media is formatted, this step is skipped.

Opening Device: Once the media is mounted and formatted, TapeWare readies the media and device for the job.

Running: After opening the device, TapeWare runs the job. As the job is run, the **Status** window automatically displays current information about the job, including which files are being backed up, restored, or verified, which streams are active, and the rate (or **throughput**) at which files are being written to media or volumes. You can use the **Display** list box to check the progress of individual streams.

Waiting for Next Pass: If some files were not backed up, restored, or verified during the first pass, such as when they are being used by other users, TapeWare attempts to access these files on subsequent passes. This message is displayed while TapeWare waits for the number of seconds specified in the **Number of passes** field on the **Options** tab of the job.

Closing Device: When TapeWare closes a device, it displays this message.

Building Logs and Audit Trails: After closing the device, TapeWare updates the storage management database with new information from the job, such as which files were backed up, and creates a log of the job.

Merging Groups: After a backup job has run, TapeWare updates the storage management database to reflect any changes to media or files created by the current job. Prior to running a restore or verify job, TapeWare sorts all the selected files into the order in which they appear on the media and displays this message.

Completed: This messages is displayed after the job is finished.


Terminated: When a job is forced to quit, TapeWare displays this message.

Viewing and Printing the Job Log

Each time a job is run, TapeWare creates a new log for that job. You can use this information to check if a job is running as you intended and to keep a permanent record of that job.

You can view the logs for a job on the job's **Logs** tab. Note that there is a separate log for each time a job is run.

The **Logs** tab shows a different log for each time a job is run.



You can specify what information TapeWare should write to the log on the **Options** tab of the job in the **Log options** field. The log always includes summary information about the job, which includes useful information about which option parameters the job used when it ran. Depending on what **Log options** parameter you selected, TapeWare will also include information about which files were successfully or unsuccessfully restored, verified, or backed up.

For further information on selecting a text editor, see "Preferences Window," Chapter 12.

To view the log of a particular job, open the **Logs** tab of the job. Select the appropriate log in the **Log Date/Time** box. When you click the **View** button, TapeWare opens the log a text editor. (You specify which text editor you want to use in the **Preferences** window.)

To print a particular log, select it and then click the **View** button. After the text editor opens print the log from the text editor. Note that some logs can be very long; check the length of the document before printing it.

Note You can also view all the available job logs from the **Instructions and Logs** option on the **Wizard** tab.

If you run a job repeatedly, you may want to delete old logs. To do so, select the job and then click the **Delete** button. Note that you can select multiple logs to delete by using the SHIFT key while selecting the logs or by dragging the mouse over multiple logs.

Note The maximum number of logs per job is 64. TapeWare will overwrite the oldest log when you reach this maximum.

Email Support for Job Logs

TapeWare provides a convenient method for you to automatically be notified if a job has run successfully or has failed to run. You can configure TapeWare to automatically email you the log of job after it has run. This way you can know for certain that a job has run successfully or the reason why a job has failed to run correctly. If you can remotely retrieve your email while you are away from the office, this allows you to monitor jobs even when you are out of the office.

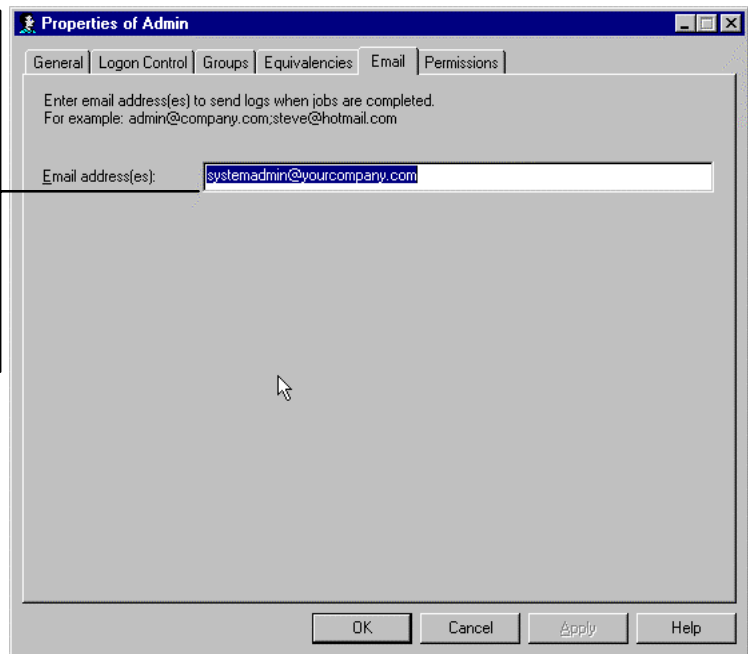
In order to use this feature, you must install the optional email support package. To install that package, start the Setup.exe program on the installation CD-ROM and select the **Install an Option**. Additional information about installing and configuring email is found in Appendix II, "Configuring Email Support."

For more information on the owner of a job, see "How Forcing Jobs to Run Effects Permissions," earlier in this chapter.

After you have installed the email support package, you must enter a valid email address for the recipient of the job logs. TapeWare will email the job log to the job's **Owner** (as listed on the **Queue** tab). Generally, if you have scheduled the job, you are the owner of the job and the job log will be sent to the address listed on the **Email** tab of the property sheet of your User object. However, if someone else is the owner of the job (because, for example, they have forced the job to run) you can still receive the job log *if your email address is listed on the **Email** tab of that user's property sheet.*

Note Be sure to configure your Email provider on the storage management server. Select **Configurations** from the **Tools** menu.

The log of a job will be emailed to the job's owner. After installing the optional email configuration package, enter a valid email address on the **Email** tab of the property sheet of all users who need to receive their job logs via email, possibly the owner of that job.



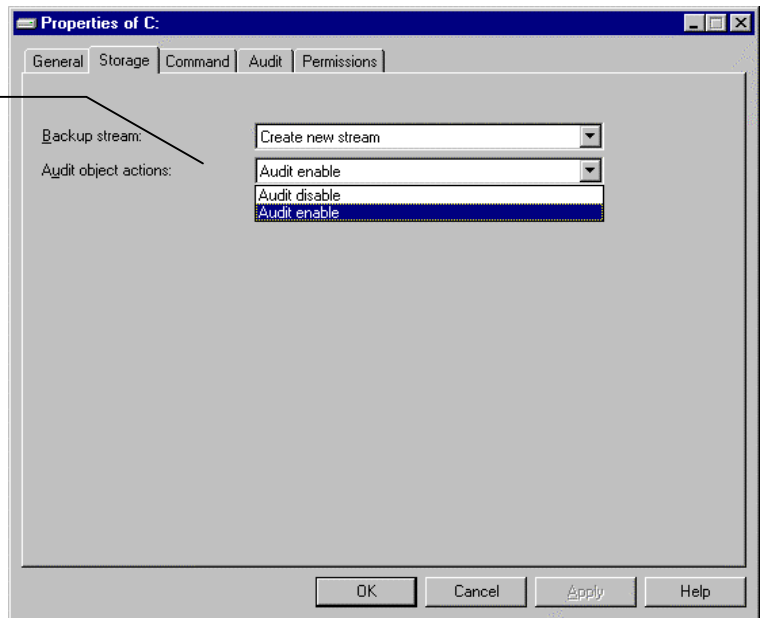
Viewing and Printing Audit Logs

Some files, such as databases, are mission critical and regular backups of these files are essential. It is also essential that system administrators be able to verify that these files have been regularly backed up. TapeWare's *audit trails* allows you to collect, store, and print such information about selected files and databases.

You can use TapeWare's auditing feature to track how often and when a file, folder, volume, or database is backed up, verified, and restored. TapeWare will create an *audit trail* for each object which is *audit enabled*. Each time an action is performed on this object, the audit trail or log is updated with data about when that object was backed up, restored, and so forth. The audit log also includes information about the media on which instances of a file are stored.

To create an audit log for an object, you must enable the audit storage property. To do this, open the property sheet for the object and click on the **Storage** tab. Change the **Audit object actions** setting to **Audit enable**. (You can only create an audit trail for objects that have a **Storage** tab, including only files, directories, and volumes.)


To create an audit log for a file, directory, or volume, set the **Audit object actions** setting on the object's **Storage** tab to **Audit enable**.



When an object has been audit enabled, a new tab appears on its property sheet, the **Audit** tab. Click on this tab to view the audit trail of that object. You can also print the audit log or save it to a file.

When an object's audit property has been enabled, the **Audit** tab appears on its property sheet.

The audit log of an object tracks when that object has been backed up and the media on which instances of it are stored.



The Media, Device, and Database Tabs

Using the **Media** tab, you can create media folders and media in the TapeWare database and delete them from the storage management database as well. The **Device** tab allows you to perform physical operations with the backup device, such as erasing, formatting, and ejecting media. The **Database** tab displays all of the objects in the TapeWare storage management database on a single tab. Many commands can also be executed from this tab.

In This Chapter

- Overview
- The Media Tab
- The Device Tab
- The Database Tab

Overview


The **Media** tab displays media folders and media objects. You can use this tab to work with these storage management database objects. For example, you can create and delete media folders, as well as create and delete media objects, such as tapes. The **Device** tab, on the other hand, is used to perform physical operations with the backup device. For example, media can be erased and formatted from the **Device** tab.

The difference between the **Media** tab and the **Device** tab is significant: the **Media** tab is used to make changes to the storage management *database* while the **Device** tab is used to perform operations using the physical devices themselves (both media and drives). When you want to make changes to the storage management database, use the **Media** tab. When you want to work with the physical media or with the device itself, use the **Device** tab. For example, if you want to change the name of a tape, you make that change on the **Media** tab because you are making a change to the TapeWare database. However, if you want to identify a tape by reading the its header, you must use the **Device** tab.

The **Database** tab displays all of the objects in the current storage management database. This may be useful, on the one hand, because you will be able to see and work with all the database objects at once. On the other hand, because all of the objects are displayed, it may be difficult to work with this tab efficiently.

The Media Tab

The **Media** tab displays media folders and media objects in the TapeWare storage management database. You can use this tab to work with these database objects. For example, you can create new media folders and media objects on this tab, as well as delete media folders and media.




The **Media** tab displays media folders and media objects.

Although you can manipulate media and media folders on the **Media** tab, many of these operations are normally handled automatically by TapeWare. For example, when you run a scheduled (rotation) job, TapeWare will automatically create any new media and media folders needed for that job. As tapes and other media are overwritten and deleted from the storage management database, these changes are automatically reflected on the **Media** tab as well.

Creating Media Folders

For further information on automatically creating media folders, see "Backup Options Automatically Updated," Chapter 7.

You might wish to create a media folder in which to store media before you create and run a backup job. Note that TapeWare automatically creates new media folders when it runs scheduled jobs. TapeWare will create a media folder in the User/Group folder and give it the name of the scheduled backup job. However, if you run a job that is not scheduled, you may want to create a new media folder in which to store the new media for that job.



Media and media folders are automatically created when an automatic rotation job is run.

◆ To Create a New Media Folder

1. Select the existing folder in which you want to store the new **Media** folder. (It cannot be the **Home** folder.)
2. Create the new folder by either
 - selecting **New Object...** from the **File** menu, or
 - clicking the right mouse button in the TapeWare object detail area and selecting **New Media Folder** from the shortcut menu.
3. Type in the name of the new folder in the tree view area.

You can also create a new folder in one step by clicking on it with the right mouse button in the tree view area and selecting **New Job Folder** from the shortcut menu. TapeWare will create the new folder inside the folder you selected.

Deleting Media Folders

When you delete a media folder, you also delete all of the objects contained within it, including any other media folders and media. You might want to delete media folders that were used by jobs you no longer plan to run and that contain media that are no longer in use.



The Delete button.

To delete media folders, select the folder and then select **Delete** from either the **Shortcut** menu or the **Edit** menu. Alternatively, you can click on the **Delete** button on the **Toolbar**.

Before deleting any media folders, you might want to move any media stored in those folders to another folder. For example, you might create a new folder named **Old Media** and move any currently unused media to this folder before deleting the media folders.

Creating New Media

TapeWare automatically creates new media when it runs backup jobs. This is true for both unscheduled and scheduled (rotation) jobs. However, there may be occasions where you may want to manually create new media prior to running the job. For example, you might be recycling old tapes you no longer use. To be certain that the tapes are labeled correctly and that no tapes are inadvertently overwritten, you might want to manually create new media before running the job.

Note, however, that for scheduled (rotation) jobs, TapeWare looks for media with specific names in specific folders. If it does not find the precise media it is looking for, the job may not run. For this reason, it is best to let TapeWare

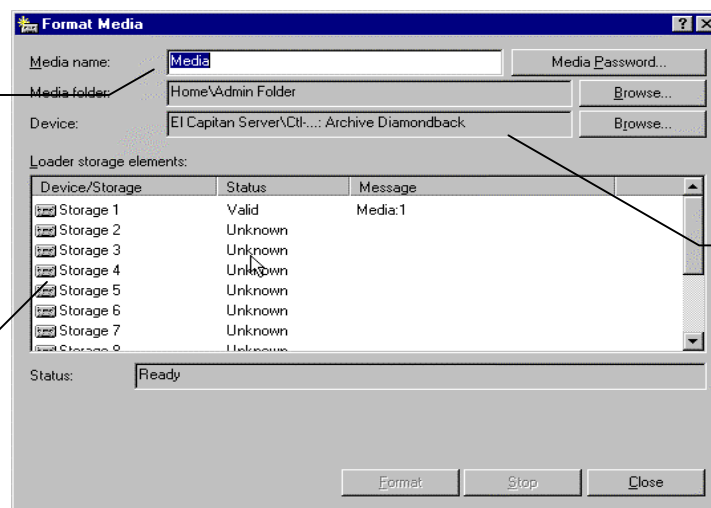
automatically create its own media for automatic rotation jobs than to manually create the media in advance.

When you create new media, TapeWare does two things: (1) it creates a new database object; and (2) it physically formats the current media in the device. This will cause any current data on that media to be lost (unusable). When you create new media, be certain that the media TapeWare will format is no longer needed.

When TapeWare formats the new media, it opens the **Format Media** window. Use this window to name the media and select a media folder in which to store the media. You must also select a device. TapeWare will format the media currently loaded in the device you select. If you select an autoloader, select the media you want to use from the storage slots that holds the media..

When you format new media, type the name of the new media in the **Media name** field.

If you are using an autoloader, select the storage slot that stores the media to be formatted.



Use the **Browse** button to select the device you want to use to format the new media.

When you format media, you can also assign the media password. See the section “Media Passwords” below for more information.

◆ To Create a New Media

1. Create the new media by either
 - selecting **New Object...** from the **File** menu and then selecting **New Media**, or
 - clicking the **New Object...** button on the **Toolbar** and then selecting **New Media**, or

- clicking the right mouse button in the TapeWare object detail area and selecting **New Media** from the shortcut menu.
2. Type in the name of the new media in the **Media name** field.
 3. If the media folder in which you want to store the media is not displayed in the **Media folder** field, click the **Browse...** button and select the proper media folder from the **Browse** window.
 4. If the device you want to use is not displayed in the **Device** field, click the **Browse...** button and select the proper device from the **Browse** window.
 5. If you are using an autoloader, select the storage slot that holds the media you want to use.
 6. If you want to assign a password to the media, click the **Media Password** button, and then enter and confirm the password in the **Media Password** window.
 7. Click the **Format** button.

Media Passwords

Whenever you format media, you can assign that media a password. By default, there is no password.

Media passwords are only required on one occasion: when media is *imported* from one storage management database to another database. For example, you might import media from an earlier version of TapeWare to the latest version of TapeWare. Or, alternatively, you might want to transfer data from one storage management zone to another storage management zone. To prevent unauthorized transferring of tapes between secure storage management zones, TapeWare lets you assign a password to any media you create. That password will be required before that media can be imported into a new storage management database.

Note Media passwords are the only security measure that prevents tapes from being imported into another storage management database. For sensitive data, be certain that every media is assigned a password.

Whether or not you assign media a password depends on your particular security needs. Media that has no password can be easily imported into any storage management database. If you do not assign the media a password, mere possession of the tape or media is enough to compromise the security of your data.

Note Previous versions of TapeWare automatically assigned media and tapes the default password “PASSWORD”. If you are attempting to import media into your storage management database from earlier versions of TapeWare and cannot do so because you have the wrong password, try using “PASSWORD” when prompted for the media password.

Deleting Media


When you delete media, TapeWare deletes information about that media from its storage management database. This includes any instances of files stored on that tape, which are also deleted from the storage management database.

Note however that deleting media does not physically erase the media. The media remains unchanged; only the database is changed. This means that you can still import that tape to another storage management database or, if you wish, back into the original storage management database.



The Delete button.

To delete media, select the media you wish to delete and then select **Delete...** from the **Shortcut Menu** or the **Edit** menu. Alternatively, you can click the **Delete** button on the **Toolbar**.




When you delete media, you only remove it from the database. Deleting media does not erase the media.

The Device Tab

You can use the **Device** tab to perform physical operations on any backup device in the current storage management zone.

Device **Controllers** are software drivers used by the storage manager. They act like containers in the database.



On the **Device** tab, **Devices** are shown as objects contained inside **Controllers**.

Any backup devices in the current storage management zone can be displayed in the object detail area. Note that this tab displays two separate types of objects, **Controllers** and **Devices**. Controllers are software drivers TapeWare uses to work with the physical devices; device objects represent the physical device itself. Each different type of device has its own unique controller. In the

TapeWare storage management database, the controllers work like a container and may have multiple devices (of the same model and manufacturer) stored within them.


When you work with devices on the **Device** tab, you must select the *device* in the object detail area, not the *controller* in the tree view area.

Note that autoloaders have two or more drivers associated with them: the **Loader** driver and one or more **Device** drivers. In general, many commands on the **Device** tab can be performed with any one of the drivers selected. (Models and manufacturers vary however.) However, if your autoloader supports multiple devices (for example, it has more than one tape read/write device) and you want to use a specific device, you must select that **Device** driver to use it. If you select the **Loader** driver, TapeWare will use the first available device in the autoloader it finds.

Restarting Failed Devices

Sometimes you will need to restart a device that has, for some reason, failed to initialize properly. A device may have stopped for any number of reasons, such as a power failure or a connecting cable malfunction.

When a device has failed, it is displayed with a yellow warning icon. The controller will sometimes display that same icon.



To restart the stopped device, you must exit or close TapeWare and then restart TapeWare. When TapeWare restarts, it will initialize the device driver again. Check the **Device** tab to see that the devices are now properly working and that they no longer display the yellow warning icon.

Device Tab Commands

After you have selected a device in the object detail area, you can perform physical operations with this device. Some of these operations effect the device itself, while others effect the current media in the device.

The following commands can be found on the **Device** menu. Many of them also have buttons on the **Toolbar** or are available from the **Shortcut** menu.

Eject Media

You can use this command to eject media from the selected device. If this command is grayed out, either your device does not support this command or no device is selected.

Eject Magazine

You can use this command to eject media magazines from the selected autoloader. If this command is grayed out, either your device does not support this command or no device is selected.

Rewind

You can use this command to manually rewind tapes in the selected device. If this command is grayed out, either your device does not support this command or no device is selected.



The Rewind button.

Retension

The **Retension** command retensions the current tape in the device by fast forwarding the tape to the end of the tape and then rewinding it to the beginning. This command can be useful in some circumstances. Occasionally when a tape is repeatedly fast forwarded and rewound for only short distances, tension differences develop in the tape that cause the tape drive to falsely believe it has reached the end or beginning of the tape. By retensioning the tape, you can sometimes make an otherwise unusable tape operational again.

If you need to retension tapes regularly in order to use them, your tape drive may need servicing, or alternatively, the tapes you are using may need to be replaced.

Clean Device

The **Clean Device...** command will run the backup device through a cleaning cycle.

For further information on setting up a cleaning cartridge on an autoloader, see "Status Tab," Chapter 12.

This command is supported only by autoloaders. If a device in an autoloader provides notification that it needs cleaning and the autoloader has a cleaning cartridge available, a cleaning cycle will be performed automatically at the start of a backup job. If you are using a device that is not an autoloader, you must manually clean the device at the manufacturers suggested intervals.

To clean an autoloader, highlight the device and select **Clean Device...** from the **Device** menu. If you are using an autoloader, TapeWare will check to see if one of the slots holds a cleaning cartridge. If it does, the cleaning cycle will be performed in the background; if not, an error message is shown.

If the **Clean Device...** command is grayed out, it is not available for your backup device. In this case, a cleaning cycle can often be performed by manually inserting a cleaning cartridge into the backup device.

Erase Quick & Secure Erase

These commands erase the media currently loaded in the selected device.



The Erase Quick button.

The **Erase Quick** command erases the first block and then writes an END OF DATA marker to that first block. The other blocks of the tape are not erased, but when that tape is read, TapeWare treats it as if it were blank because it encounters the END OF DATA marker in the first block.



The Secure Erase button.

The **Secure Erase** command erases every block of the tape. This operation can be very time consuming, lasting up to two hours. However, it will physically erase every block of the tape. If you are attempting to destroy sensitive data, you may want to use the **Secure Erase** command.

Not all devices support both commands. Many devices support only one of the two erase commands. If one of the commands is grayed out, the selected device does not support that command. If both commands are grayed out, no device is currently selected.

Format Media

You can use this command to format media currently loaded in the selected device.



The Format button.

When TapeWare formats new media, it opens the **Format Media** window. Use this window to name the media and select a media folder in which to store the media. You must also select a device. TapeWare will format the media currently loaded in the device you select. If you select an autoloader, select the storage slot that holds the media you want to use.

When you format new media, type the name of the new media in the **Media name** field.

If you are using an autoloader, select the storage slot that stores the media to be formatted.

Device/Storage	Status	Message
Storage 1	Valid	Media:1
Storage 2	Unknown	
Storage 3	Unknown	
Storage 4	Unknown	
Storage 5	Unknown	
Storage 6	Unknown	
Storage 7	Unknown	
Storage 8	Unknown	
Storage 9	Unknown	

Use the **Browse** button to select the device you want to use to format the new media.

When you format media, you can also assign the media a password. See the section “Media Passwords” above for more information.

◆ To Format Media

1. Select **Format Media** from the **Device** menu or the **Shortcut** menu. Alternatively, click the **Format** button on the **Toolbar**.
2. Type in the name of the new media in the **Media name** field.
3. If the media folder in which you want to store the media is not displayed in the **Media folder** field, click the **Browse...** button and select the proper media folder from the **Browse** window.
4. If the device you want to use is not displayed in the **Device** field, click the **Browse...** button and select the proper device from the **Browse** window.
5. If you are using an autoloader select the storage slot that holds the media you want to format.
6. If you want to assign a password to the media, click the **Media Password** button, and then enter and confirm the password in the **Media Password** window.
7. Click the **Format** button.



The Identify Media button.

Identify Media

You can use this command to get the name of the media currently loaded in the device. When issued this command, TapeWare checks to see if it knows what tape or other media is currently loaded in the device. If it doesn't, TapeWare reads the header of the current media, a process that may take up to several minutes. When TapeWare has learned the name of the media currently loaded, it displays that name in the **Identify Media** window.

Import Media

This command allows you to use data on media that was created in another storage management zones or by another software program. To use media that was not created in the current storage management database, you must import that media into the current database.

You might import media in one of four situations:

- when you want to use media created by an earlier version of TapeWare.
- when you want to use media created in a different TapeWare storage management zone.
- when you want to use media created by another backup program.
- when you want to use media accidentally deleted from the storage management database.

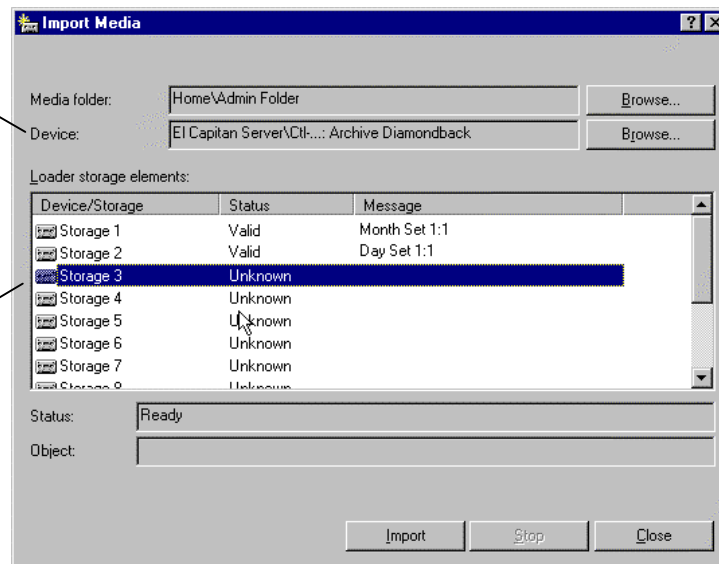
TapeWare can import media created by earlier versions of TapeWare, from TapeWare 3.X and above. You can also import media created by ARCserve, from versions 4.X and above. If you import media created by ARCserve, that media must have been created in non-SMS format, that is, it cannot have been created using Novell's proprietary SMS format.

When you import media, you must supply the media password. No password is required if the media has no password. Note that earlier versions of TapeWare assigned media the password "PASSWORD". If you are having trouble importing a tape created with an earlier version of TapeWare, try using "PASSWORD" as its password.

TapeWare will not perform any other operations while it is importing media. Additionally, the process may last up to two hours. Before you import media, be certain that there is sufficient time to complete this lengthy process. Additionally, you will want to be available to log out of TapeWare when the import is complete, in order that the security of the LAN is not compromised.

When you import media, select the device you want to use to read the imported media.

If you are using an autoloader, select the storage slots that holds the media you want to import.



◆ To Import Media

1. Select **Import Media** from the **Device** menu or the **Shortcut** menu.
2. The **Media Folder** field displays the folder in which TapeWare will store the imported media. If this is not the folder in which you want to store the media, click the **Browse...** button and select the proper media folder from the **Browse** window.
3. The **Device** field displays the device from which TapeWare will import the media. If the device you want to use is not displayed, click the **Browse...** button and select the proper device from the **Browse** window.
4. If you are using an autoloader, select the storage slot that holds the media you want to import.
5. Click the **Import** button.
6. In the **Import Password** window, enter the password for the media you are importing and then click **Ok**. If the media has no password, leave the field blank and click **Ok**.

Restore Database

This command provides a quick method of restoring your current storage management database in case it has been corrupted. Use **Restore Database...** command when your current set of media is intact, but the storage management

database has been lost or corrupted for some reason. For example, you might use this command if the storage management server has crashed.

The **Restore Database...** command differs significantly from the **Import Media...** command. Use the **Import Media...** command when you want to add media to the current storage management database, such as when you are adding tapes from another TapeWare zone. The **Import Media...** command *does not replace* the current storage management database; it only adds additional data to it. The **Restore Database...** command, on the other hand, *replaces* the current storage management database with the last known good database. The advantage of the **Restore Database...** command is that it provides a quick and easy way of replacing a lost or corrupted TapeWare storage management database. (You could use the **Import Media...** command to restore a corrupted database, but this process is very time consuming and, if you have multiple tapes, might require many hours or even days.)

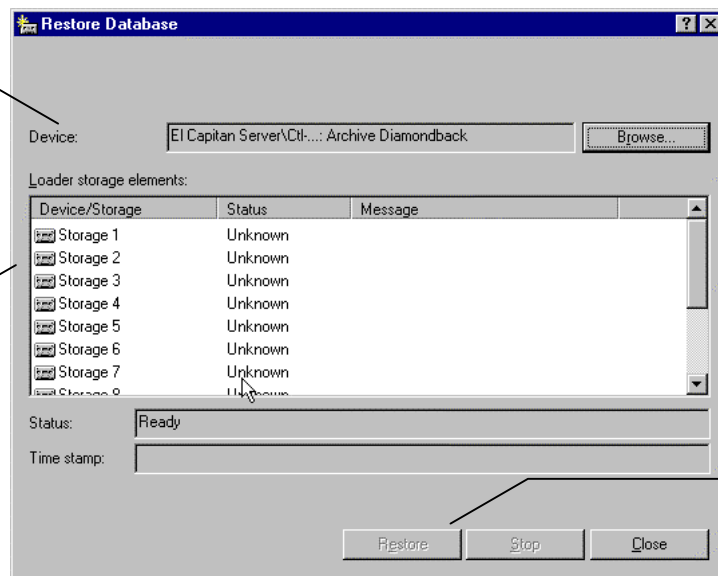
To use this command, locate the last media on which the storage management database was backed up. If you have printed the job log from the last backup job, you can read the job log to see on which media the database is located. (To make certain the database is regularly backed up, check the **Selection** tab of the backup job to see that the database for your zone is checked.) Insert the media into the backup device, select the device in the object detail area of the device tab, and then select **Restore Database...** from the **Device** menu. In the **Restore Database** window, check to see that the proper device is shown in the **Device** field. If not, click the **Browse...** button and select the correct device. If you are using an autoloader, select the correct storage slot in the **Device/Storage** field. Click **Restore**.

After the command is complete, in order for the restored storage management database to be used, you must exit or close TapeWare and then restart it. When TapeWare restarts, it updates the database to match the restored version.

If TapeWare is running as a service, you must stop, then restart the service. You may use the TapeWare Service Control Manager to start and stop the TapeWare service.

To restore a corrupted or lost database, select the device you wish to use on the **Device** tab and select **Restore Database...** from the **Device** menu.

If you are using an autoloader, select the proper storage slot...



...and then click **Restore**. After the command completes, quit and restart.

Note All current information in the current TapeWare storage management database will be lost when you use the **Restore Database...** command. This command *does not append* data to the current storage management database—it replaces the current storage management database with the last known good database. Use this command only when you want to restore a corrupted or lost database.

◆ To Restore a Corrupted or Lost Storage Management Database

1. Locate the media on which you have backed up the storage management database you wish to restore. Normally, this is the last backup job run.
2. Insert the media into the appropriate backup device and then select it in the object detail area of the **Device** tab.
3. Select **Restore Database...** from the **Device** menu.
4. If you are using an autoloader, select the proper storage slot from the **Device/Storage** field.
5. Click **Restore**.
6. After the restore is complete, exit or quit TapeWare. When you restart TapeWare, the database will be restored.

Note You *must* quit or exit TapeWare before the storage management database will be restored. TapeWare only completes the restoration of the database when it is restarted.

The Database Tab

All of the objects in the current TapeWare storage management zone are displayed on the **Database** tab. (Like other tabs, of course, you will only be able to see those objects to which you have permissions.)

Many commands can be performed from the **Database** tab, including all of the commands on the **Device** tab and the **Queue** tab. These commands include running jobs, stopping jobs, formatting media, and so forth. Note however that you cannot create new objects on the **Database** tab.

The advantage of the **Database** tab is that it can display all of the objects in the storage management database at once. This is also its disadvantage, that so many objects are displayed that it may be cluttered and difficult to work with.



The Query button

You can use the **Query** window, however, to limit what files and objects are displayed on this tab. The **Query** window works similar to the **Selection Filters** window for jobs. It sorts through database objects and only displays those which meet the specified criteria.

You can use the **Query** window to 'sort' through the objects in the database in order to only display those objects which match the criteria you have specified.

For example, you might want to only display files that have not been backed up or only those files that have been deleted. By setting the appropriate filters, you can control which objects are displayed on the **Database** tab. For additional information on how each of the query filters work, see “Query window” in Chapter 12.

Tips, Techniques, and Strategies

This chapter contains information you can use to work more efficiently with TapeWare. The first sections explore managing the storage management database and running jobs faster. The last sections of the chapter cover practical techniques for working with jobs.

In This Chapter

- Managing the Storage Management Database
- Strategies for Faster Jobs
- Working with Permissions
- Working with Scheduled Jobs
- Selecting Files for Jobs
- Restoring Tips
- Other Tips

Managing the Storage Management Database

An important decision when planning a comprehensive backup strategy is where to locate the TapeWare storage management database. This section explores some considerations you should review before making this decision.

Where to Locate the Storage Management Database


- *Consider locating the database on a workstation or file server other than the main file server.*

Recall that TapeWare keeps track of objects and properties in a storage management database that it creates and manages. Where should you store this database? That is, on what volume and machine should be storage management server?

TapeWare lets you locate the storage management database on any machine (workstation or file server) or volume in the storage management zone. Note that the storage management database does *not* have to be on the same machine as


the backup device. It could be placed on any volume of any machine that is a member of the storage management zone.

You can locate the storage management database on any machine in the TapeWare network, including, in this case, a workstation. Attach the backup devices to machines using local buses for greatest speed.



For example, suppose you have a file server with a large RAID device attached. Backup jobs using this RAID device will run fastest when the backup device is placed on the same machine as the RAID device. On the other hand, the storage management database may be best located on another machine other than the file server. This is because if the file server were to become inoperable (for example, the drive were to crash), you would still be able to use the storage management database to restore the file server volumes. Had the storage management database had been located on the file server, however, the database must first be restored before other files could be restored. This can be a lengthy, time-consuming process.

A good strategy to consider is to place the backup device on the file server for maximum speed, but to locate the storage management database on a separate machine. Consider this example. Two file servers are connected on an Ethernet network. Each file server has its own backup device, which helps the jobs run faster and more efficiently. The storage management server for all three machines is located on a separate machine (which can be called a “storage management server”). Jobs can be run from this storage management server; additionally, they can also be run from any other machine on the network as well.



(Note that the backup device could just as easily be placed on any other machine in the network. TapeWare does not require that the backup device be physically attached to a file server. Additionally, while the term “storage management server” is a convenient label, in fact, TapeWare does not require that this machine be a network server. It could just as well be a client machine or workstation.)

This arrangement has several features to recommend it. First, jobs run quickly because most of the data is transferred over local buses, instead of over the network. In this arrangement, TapeWare will automatically route data from each of the file servers over local buses to its own backup device. Whenever there is a choice, TapeWare automatically routes data over local connections rather than network connections.

Second, administering jobs is uncomplicated. Jobs can be created and run from any machine on the network. Note that jobs can also be administered from a machine that is a member of a *different* storage management zone. The TapeWare administrator or other user can log on to this storage management database from another database zone and then create and run jobs in that zone.

Third, suppose a disaster occurs and the RAID device of one of the file servers needs to be replaced. Because the storage management database is located on another machine, recovery is quick and easy. The storage management database contains all of the information necessary to restore the lost data. Had the database been stored on the file server, recovery would have been much more difficult. Note that the backup device on the other file server can be used to help restore the file server that failed.

For further information on restoring a corrupted storage management database, see "Restore Database..." Chapter 9.

Fourth, suppose the storage management server fails and the TapeWare storage management database is lost. While the loss of the storage management server is significant, the file servers are not impacted and they can continue to perform their tasks. Backups can still be run from another machine on the network, or if need be, from the file servers themselves. Since the storage management database was lost, it will have to be imported from existing tapes. This step, however, is not crucial and can be run when convenient, as long as careful backup procedures are followed (that is, no important tapes are overwritten and a full backup job is run immediately).

The advantages of this arrangement extend to any backup strategy or network arrangement. In general, consider placing the TapeWare storage management database on a machine separate from the most important data. At the very least, consider locating the database on a separate volume.

Accounting for Storage Management Database Size

Before deciding where to locate the database, consider how large the TapeWare database might eventually become.

The size of the storage management database is primarily a function of the number of tapes in the backup set, the number of files backed up, and the number of instances of each file on valid media. To a lesser extent, the number of objects in a storage management database and the properties of those objects affect the size of the database.

In its storage management database, TapeWare assigns

- 40 bytes for each instance of a file,
- 128 bytes for each database object, and
- 1024 bytes for the properties associated with objects.

You can use these figures to estimate the size of the storage management database. Files are by far the most numerous of the objects in the database, so much so that the size of the other objects (such as users and jobs) is negligible. Each file that is backed up is made part of the storage management database and is assigned 128 bytes. While this figure can be significant, the *number of instances* for that file in the storage management database are more important when predicting the size of the database. This is because there might be as many as 20 or 30 instances of a particular file in the storage management database. (Files only have properties associated with them when they are assigned permissions or storage streams, but generally, since most files inherit their permissions from other objects, the 1024 bytes assigned for properties can be ignored.)

Consider the following example. Suppose a file server which holds 100,000 files is regularly backed up using the GFS-25 automatic rotation schedule. Because there are 25 tapes in this backup set, there are potentially 25 instances of each of these files (although in practice, there will be fewer instances since most of the jobs are incremental). You can predict the size of the storage management database as follows:

$$[(\# \text{ of files}) * (128 \text{ bytes})] + [(\# \text{ of instances}) * (\# \text{ of files}) * (40 \text{ bytes})] ; \text{ or}$$
$$[(100,000) * (128 \text{ bytes})] + [(25) * (100,000) * (40 \text{ bytes})] = \sim 108 \text{ MB.}$$

Storage Management Server and Machine Platforms

Finally, consider the operating system of the machine of the storage management server. For multi-user, LAN installations, place the database on either a Windows NT server or a NetWare server. While the storage management database can be located on a machine running Windows 95/98, this option is not recommended. Performance on network installations will be increased when using an NT or NetWare server.

For single machine installations, TapeWare also runs on Windows 95/98 and DOS platforms, although TapeWare runs more efficiently with Windows 95/98 than with DOS.

Strategies for Faster Jobs

TapeWare is designed to run jobs quickly and efficiently on various network platforms and arrangements. This section will help you plan your backup strategy and network installation to maximize TapeWare's speed and efficiency.

What Slows Down Tape Drives?

Tape drives work most efficiently when they are "streaming," that is, when the tape itself is constantly moving forward and a steady flow of data is being written to the tape as it passes the write head. In general, tape drives will "stream" if there is a constant flow of data available to write to the tape.

Whenever there is an interruption in the data and the drive must wait for the data, the "stream" breaks off and the tape stops. Additionally, after the tape stops, the drive must reposition the write head and regain the velocity of the tape. To do so, the drive will back up the tape prior to where it stopped writing, and then restart the tape forward again. This can be a time consuming process, especially if it occurs repeatedly.

For this reason, it is important to keep data constantly flowing to the tape drive. The performance of your job will be maximized when the tape drive constantly has data available to it.

Maintaining the Flow of Data

There are several steps you can take to make sure that data is constantly available to the backup device.

Prefer Local Busses to Network Connections In general, data is transmitted more efficiently over a local bus than over a network connection. Thus, for example, performance will increase if the backup device is placed on the network file server, rather than on a client machine, although either arrangement is possible.

Note that the storage management database need not be placed on the same machine as the backup device. The location of the storage management database does not affect the rate of data transmission. (Other features of the jobs, such as building backup selection lists and opening and closing files, create network traffic and require CPU calculations. While the location of the storage management database will affect these parts of the job, the advantages of a remote location of the database often outweigh any speed loss.)

Add a Second Device Controller Even though a backup device and volume may be on the same machine with a local connection between them, if the backup device and the volume share the same device controller, this limits the speed of data transmission. Adding a second device controller can effectively double the rate of data transmission. For example, consider adding a second SCSI controller for the exclusive use of the backup device. This can result in a dramatic improvement in the rate of data transmission.

Use a Wider Data Stream For both network and local connections, the width of the data stream effects the rate at which data can be transferred. Many tape drives are capable of writing data much faster than can be transferred over older device controllers and network connections. By upgrading to more recent technology with wider data streams, there will be fewer interruptions in the stream of data to the backup device. For example, consider using a Wide or Ultra Wide SCSI controller for the backup device. Additionally, for network connections, consider upgrading from 10Base-T Ethernet to 100Base-TX.

Consult your manufacturer's documentation to see if your backup device would benefit from a wider data stream.

Use the Proper Number of Data Streams TapeWare is capable of controlling up to 8 data streams per backup device simultaneously, up to a maximum of 16 streams. This feature can greatly increase your job speed since multiple devices can simultaneously send data.

You control the data streams of a volume, directory, or file with its **Storage** tab. Normally, the **Backup stream** field on the **Storage** tab of a *volume* is set to “Create new stream,” while the **Backup stream** of a *directory* or file is set to “Use existing stream.” By changing these options, you can change the number of backup streams, thereby effecting the speed at which the job runs. Assigning the proper number of data streams can help speed up the data transmission rate.


In general, you should assign data streams according to the number of streams the *physical* device (e.g., the disk drive) is capable of handling. Usually this number is equal to the number of spindles the drive has. Under most circumstances, you should use this number to determine the number of data streams.

There is one exception to this general rule, however. When working with very large files, performance may be increased by creating an additional stream for these very large files. In particular, if you have a very large file, .5 GB or larger, when you create a new stream for this file, TapeWare may be able to send data to the tape drive at a rate that will allow for “streaming.”

Consider these three examples. (1) A RAID device is capable of sustaining multiple data streams at once. By assigning various directories on the RAID device additional streams, you can increase the rate of data transmission. To do so, assign large directories their own stream by changing the **Backup stream** field of each directory’s **Storage** tab to “Create new stream.” Don’t add any more streams than the RAID device is capable of sustaining.

(2) On the other hand, a single physical device may have multiple logical volumes. If each volume is assigned a separate stream, this will not result in faster data transmission, and may in fact result in slower data transmission if it creates additional seeks by the disk drive. To turn off one of the data streams, change the **Backup stream** field on the volume’s **Storage** tab to “Use existing stream.”

(3) Some file servers may have large database files on them, perhaps 1 GB or larger. These files should be assigned their own stream. To do so, change the **Backup stream** field on the file’s **Storage** tab to “Create new stream.”



Use the storage tab to create new streams for large files and directories on RAID devices.

In general, when creating or modifying data streams, first use the capacity of the *physical* device to determine optimum number of data streams and then second

create separate streams for very large files. Too few or too many data streams will impede maximum performance.

Other Factors that Effect Job Speed

Additional factors effect job speed, although less dramatically.

File Compression Whether or not files are compressed by the backup unit effects how fast jobs run. When a backup device compresses files, often at a 1.8:1 ratio, this means that a proportionate greater amount of data needs to be sent to the backup device in order for it to stream. However when files are sent across the LAN already compressed, such as in the case of NetWare, further compression by the backup device will be negligible.


Minimize Small File Size Large files transfer and are written to the backup device more efficiently than are small files. If you can limit the number of small files you backup, especially those smaller than 64K, your job will run more quickly.

CPU Speed In general, a faster CPU results in faster backups. Take the speed of the CPU into account when deciding where to place the backup device and to locate the storage management database.

Turn on “Native data streams format” By default, TapeWare stores files on tape in the format that they were transmitted across the network, that is, in a format specific to either NetWare or Windows NT. In general, jobs run faster when this default setting is left unchanged.

The **Native data streams format** option is in the **Advanced options** window. In general, it should be left with its default setting (checked). However, if you plan to share data from one network platform to another, you should clear this option.

For faster jobs,
keep the **Native
data streams
format** checked...



Working with Permissions

This section provides useful tips for assigning permissions.

Checking the Effective Permissions of a User

- *Log on as the user.*

On complex installations with multiple users and groups and varying levels of security, a particular user's effective permissions can be difficult to intuit.

There is a simple way to check what effective permissions a particular user has: by logging on as the user.

If you have not yet assigned the user a password, simply log on as the user. Browse the various **General** tabs of the objects in the storage management database and check to see if the effective permissions displayed match your intended security measures.

If the user has a password and you do not know it, create another "alias" user and make that user equivalent to the user whose permissions you wish to check. Then log on as the alias user. Be certain to delete the alias user after browsing the **General** tabs of various database objects.

Using Groups to Handle Complex Security Needs

- *Set up groups and then make users members of them.*

Some security arrangements can be very complex, with multiple users possessing differing levels of effective permissions to different storage management database objects. Setting up each user's permissions separately and individually can be a complex and time consuming process.

You can use groups to speed up this process. Consider the following simplified example. Suppose you want some users to have full permissions to a tape drive (that is, the ability to create new tapes, to overwrite old tapes, to write backup tapes, and read tapes for restore jobs), but want other users to have limited permissions to the tape drive, for example, only the ability to write to backup tapes, but not overwrite them.

Begin by creating two new groups. Name one user "Users with Full Permissions to Tape Drive" and assign this group **Create, Modify, Delete, Write** and **Read** permissions to the tape drive. Name the other user "Users with Write Permission to Tape Drive" and assign this user **Write** permission to the tape drive. Next, delete the corresponding User/Group folders that appear on the job tabs.

Then when you create new users, rather than individually assigning each user permissions to the tape drive, make them members of the appropriate group.

You create as many groups as necessary, with varying levels of access to storage management database objects such as media, machines, volumes, and directories. For example, you might create a group named "Backup Permission

to Volume” and another named “Backup and Restore Permission to Volume,” assigning each the appropriate permissions.

Working with Scheduled Jobs

This section provides useful tips for running scheduled jobs.

Running a Failed Rotation Job Again

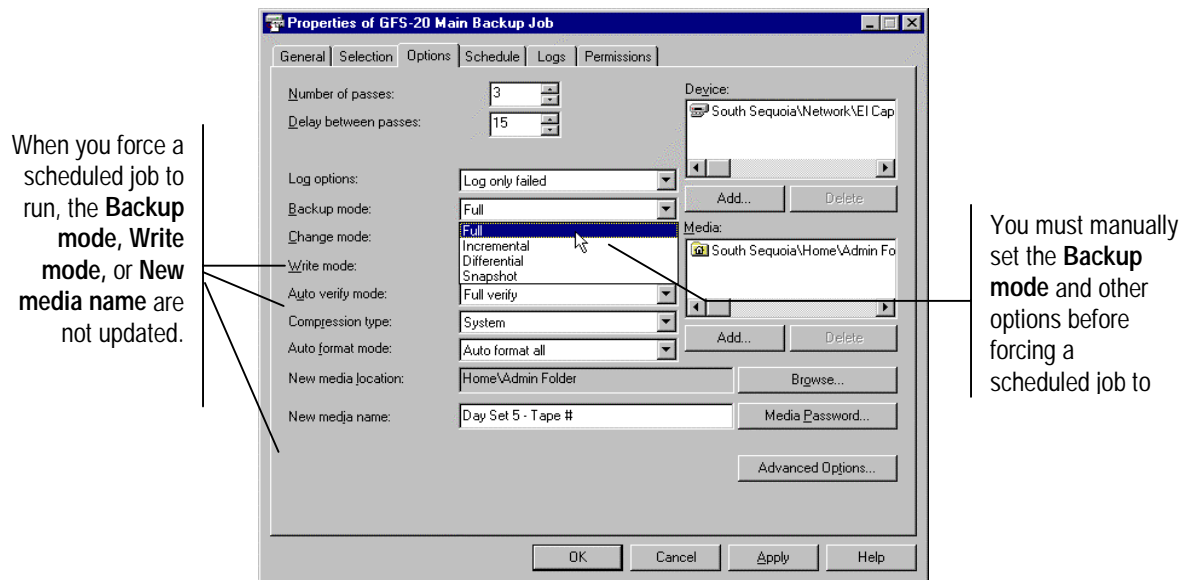
- *Manually set the correct options and “force” the job to run again.*

Suppose a scheduled job has failed to run correctly. In order to insure the integrity of the data, the job must be run again.

Consider this example. Suppose it is discovered on Monday morning that a full backup job has failed to run as scheduled on Friday evening. If a full backup job is not run before the next incremental job, the ability to fully reconstruct data will be compromised. It is vital that the full backup job be run soon.

However, you cannot merely “force” the job to run again. Recall that when TapeWare runs a scheduled job, it automatically updates three parameters on the **Options** tab of the job: **Backup mode**, **Write mode**, and **New media name**.

Note that TapeWare does NOT automatically update these fields when you manually “force” a scheduled job to run. For example, when TapeWare automatically runs a scheduled backup job on a Monday, it changes (updates) the **Backup Mode** from **Full** to **Incremental**. But when this job is “forced” to run before its scheduled time, TapeWare does not automatically update these fields.



Before forcing a failed job to run again, open the job log of the failed job, noting the appropriate options. If needed, print the job log. Next, open the **Options** tab of the failed job. Set the parameters on the options tab to match those of the failed job. In particular, check the **Backup mode**, **Write mode**, and **New media name**.

You will also want to select the appropriate media in the **Media** field. Use the **Browse** button to select the same media as the failed job was to use. When the options of the job match the options the failed job would have used, run the job.

If you have changed the **Media** field, be certain that this field is set back to its original specification so that scheduled jobs will automatically select the proper media.

An alternative method would be to copy the failed job, change the schedule type to **Not Scheduled**, and then set the options parameters to match the failed job. Manually “force” the job to run and delete it after it has successfully completed.

Scheduling a Job to Run Once Only

- Turn off all the dates on the **Manual** schedule except the desired date.

Often, when you want a job, such as a restore job, to run only once, the job is unscheduled and it runs only when you instruct TapeWare to run the job. What if, however, you want a job to run only once, but need to schedule it to run during non-peak hours?

Try this method. Begin by selecting the **Manual** schedule. Then turn off all of the days of the week by clicking on the day of week name tiles. The calendar will be all white. Then click on the day you want to run the job with the right mouse button and select **Daily**. This will be the only day the job runs. Be certain to adjust the time you want the job to run in the **Start Time** box.

Selecting Files for Jobs

This section provides useful tips for selecting files.

Selecting Files not Previously Backed Up

- Set the **Instance Range** filter to “*At most 0*”

Suppose you want to run a backup job that only selects files that have not been backed up previously. You can use the **Instance Range** filter to “filter out” any files that have been previously backed up.

Each time TapeWare backs up a file, it creates a new *instance* of that file. If a file has not been backed up, TapeWare has no instances recorded in the database for that file.



Filter Change
button

To select only files that have not been previously backed up, begin by clicking the **Selection Filters** button on the toolbar of the **Selection** tab. The **Selection Filters** window will appear. Click on the **Filter Change** button next to the **Instance Range** field. Set the **Range Type** to **At most** and then set the **Maximum instances** field to **0**. TapeWare will only select those files with no instances.

To select only files that have not been previously backed up, set the **Instance Range** filter to **At most 0**.

The screenshot shows a dialog box titled "Instances Range". It has a "Range type:" dropdown menu set to "At most". Below it are two numeric input fields: "Minimum instances:" set to 0 and "Maximum instances:" set to 0. At the bottom right are "OK" and "Cancel" buttons. A line from the text on the left points to the "Maximum instances:" field.

Note that this method does not insure that you have the latest instance of every file. Having an instance of a file does not insure that the instance you have reflects the latest changes to the file. It may have been modified after the last time you backed it up and so your latest instance may not match the file's current form.

Selecting Deleted Files for Restoring

- Set the **Delete range** filter to **On or before** some random future date.


When a file has been deleted from a file server or workstation and an instance of that file exists on valid media, TapeWare marks that file in its storage management database as having been deleted and assigns it a delete date. Additionally, these files appear with a special icon in the object detail area of the **Selection** tab.



Filter Change button

You can use the **Delete range** filter to select only those files for restoring that have been deleted. Begin by clicking the **Selection Filters** button on the toolbar of the restore job's **Selection** tab. The **Selection Filters** window will appear. Click on the **Filter Change** button next to the **Delete range** field and then select **On or before** from the **Range type** list box. Next select a random future date, for example, Jan. 1, 2020. TapeWare will exclude all of the files that have not been deleted from the set of files to be restored. When you return to the **Selection** tab, all and only those files that have been deleted will be checked.

To select files that have been deleted, set the **Delete range** filter to **On or before** some future date.



Selecting Instances from a Specific Job

- Select the appropriate instance date for a container object.

When a file is backed up, TapeWare creates an instance. Each instance of a file has a unique instance date and every file backed up during the same job has the same instance date. (You can view this information for all of the available instances in the **Instances...** window.)



The Select Instance button

Remember that when you specify the instance date for a container, such as a volume or folder, objects in that container are only selected when they have the same instance date.

If you want to select only those files backed up during a particular job, begin by checking a container high in the tree hierarchy, such as the machine or network icon. This will cause all of the files below this object to be initially selected. Then open the **Instances...** window by clicking on the **Select Instance** button. Select the appropriate date and time instance for the job. Now only those files with a matching instance date will be selected.

Selecting Instances from Specific Media

- *Add the media to the **Media** filter.*

Suppose you want to restore only those files that appear on specific media or want only to verify files from specific media. You can use the **Media** filter on the **Selection Filters** window to only select files that have valid instances on the media you specify.

To do so, open the **Selection Filters** window by clicking on the **Selection Filters** button on the toolbar of the **Selection** tab of the job. Then click the **Add...** button to open the **Browse** window. When you add media to the **Media** field, TapeWare checks to see if the file selected has a valid instance on that **Media**. If so, that file is included in the job. (If you add multiple media to the **Media** field, only files with instances on all the selected media will be included in the job.)

Restoring Tips

This section provides tips for restoring files and volumes.

Restoring Volumes for the Latest Date

If you have used a built-in schedule and have run backup jobs as scheduled, you can easily and simply restore files as they appeared the last time a backup job was run. Simply select the volumes or files you wish to restore on the **Selection** tab of a restore job. The **<Latest>** instance of each file will be automatically restored to the volume. TapeWare will prompt you for whatever tapes are needed to complete the restore job.

Restoring Volumes for a Specific Date

You can restore volumes and directories as they appeared on a particular date as long as that date is within the full data recovery period. Recall that the full data

recovery period is the number of days prior to the data loss for which any and every file backed up can be recovered. (If you want to restore for the last date files were backed up, see the section above for a simple method.)

Different schedules provide full data recovery periods for varying numbers of days prior to the last backup. For example, a GFS 30-tape job can reconstruct the data for any day of the past three weeks, while a simple 4-tape backup only provides for reconstruction of the past two days.

You can reconstruct the data for any particular day during the full data recovery period. Consider this example. Suppose you want to restore a particular volume as it appeared on Wednesday morning. Providing the date falls within the full data recovery period, there are three possible scenarios for restoring the volume as it appeared at the beginning of business on Wednesday: either (A) restore from a full backup tape; (B) restore from a full backup tape and the most recent *differential* tape; or (C) restore from a full backup tape and all of the *incremental* tapes from the previous full backup and the date in question.

Example A assumes a full backup job was run the Tuesday evening before the Wednesday for which you want to restore data.

Example C assumes incremental jobs were run on the Monday and Tuesday before the Wednesday and a full backup job on the previous Friday.

Sun	Mon	Tue	Wed	Thu	Fri	Sat
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31				

Example B assumes a differential job was run on the Tuesday before the Wednesday and a full backup job on the previous Friday.



Filter Change button

- (A) If you ran a full backup job the previous evening, you can run one restore job. Begin by creating a new restore job and selecting the appropriate volume on the selection tab of the restore job. Initially, the **<Latest>** instance of these files will be selected. You must select the instances according to the desired date, in this case the Tuesday before the Wednesday. To do so, set the **Backup range** filter to the desired date. Begin by clicking the **Selection Filters** button on the toolbar of the restore job's **Selection** tab. The **Selection Filters** window will appear. Then click on the **Filter Change** button next to the **Backup Range** field. Specify Tuesday's date in the **Date Range** window.
- (B) If you ran a differential job the previous evening, you need only run two restore jobs. The first restore job must restore all of the files from the previous full backup job; the second restore job must restore files from the previous evening's differential job.

Suppose in this example that the last full backup was done on Friday evening and that a differential job was run on Tuesday evening. To restore files as they appeared on Wednesday morning, follow these steps.

First, create a restore job, select the appropriate volume, and then set the **Backup range** on the **Selection Filters** window to match Friday's date. Name the job with an appropriately identifying name, such as, "Restore from Friday's Full Backup."

Second, copy the first restore job, rename it with an identifying name, and change the **Backup range** date to match Tuesday's date.

Run the two jobs, being certain to run them in the proper order.

- (C) *If you ran an incremental job the previous evening, you will need to run two or more restore jobs. The first job must restore all of the files from the previous full backup job; the other jobs must restore all of the files from all of the previous incremental jobs between the full backup and the date in question.*

Suppose in this example that the last full backup was done on Friday evening and that incremental jobs were run on Monday and Tuesday evenings. To restore the volume as it appeared on Wednesday morning, follow these steps.

First, create a restore job, give it an identifying name, select the appropriate volume, and then set the **Backup range** to Friday's date.

Second, copy the first restore job, rename it with an identifying name, and change the **Backup range date** to match *Monday's* date. Repeat this step, changing the **Backup range date** of this third job to *Tuesday's* date.

Run the three jobs, being certain to run them in the correct order.

Copying a Directory Structure

- *Clear the **Children** check box in the **Selection Filters** window.*

Suppose you have set up a complex directory that you want to replicate in a new location, for example, on a new workstation or file server. TapeWare provides you an easy way to do this.

If you have not previously backed up the directory, create a backup job that does so. Select the appropriate volume. Then, open the **Selection Filters** window. Clear the **Children** check box, being certain that the **Parents** box is checked. (The job will run faster when the **Children** box is cleared, however it is not necessary to clear this option. You can restore the directory by itself, even when you have previously backed up both the directory and the files in it.)

To copy the directory structure to a new location, create a restore job, selecting the appropriate directory and restore location. Then, open the **Selection Filters** window. Clear the **Children** check box, being certain that the **Parents** box is checked. The job will “copy” that directory to the new location.

Restoring Files to a New or Different Folder

- *On the Selection tab, drag the files to a different folder.*

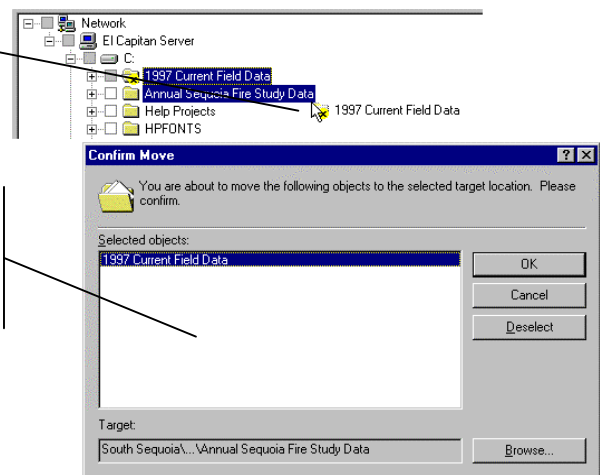
Suppose you want to restore files or folders, but do not want to overwrite files and folders currently existing on the volume. To avoid overwriting (replacing) current files or folders with the instances you are restoring, restore the files or folders to a new or different location.

When you instruct TapeWare to restore files and folders in new locations, TapeWare creates new files and folders in the specified location.

To restore a file to a different folder, drag the file in the tree view area on the **Selection** tab of the restore job to the new folder. Alternatively, if the target location is not displayed, highlight the file and select **Move...** from the **Shortcut** menu. In the **Confirm Move** window, select a target location. TapeWare will move the file to the location you specify in the **Target** field.

To restore a folder or file to a different folder, select it, drag it to the new folder...

...and then confirm the move in the **Confirm Move** window.



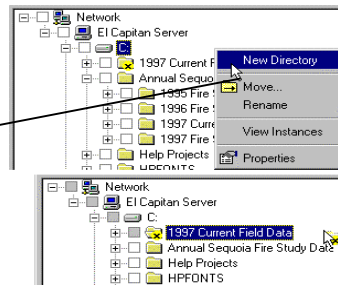
You can also restore folders and volumes in new locations. The contents of these containers move with them and are restored, along with the folder or volume, in the new location.

Additionally, you can create a new folder and restore files to that new folder. When TapeWare restores the files, it creates the new folder and restores the files

you specified to that new location. Similarly, you can restore folders and their contents in new folders you create.

To create a new folder, in which to restore the file or folder, first highlight the location where you want to create the new folder in the tree view area. Then click on the **New Object** button on the **Selection** tab's **Toolbar**. Or use the **Shortcut** menu and select **New Directory**. TapeWare will create the new folder in the location you specified. Give the folder a new name and then drag into that folder the files and folders you want restored in it.

To restore a file or folder to a new folder, first highlight the container in which you want to create the new folder and then, clicking the right mouse button, select **New Directory** from the **Shortcut** menu.



Next, drag the folder or file you want to restore to the new folder you created.

Note that when you move an instance on the **Selection** tab of a restore job, the changes you make are only reflected in the current restore job. Only the current restore job will assign the file or folder the new location. When you create a new restore job, you will see the files and folders in their original locations. Likewise, the **Database** tab will continue to display files in their original locations.


Restoring Files with New Names

- *Rename the file on the **Selection** tab of the restore job.*

Suppose you want to restore a file with a different name. To do so, you rename the file after you select it. When you rename the file, TapeWare restores the file with the new name. This can be useful for not overwriting versions of the file that currently exist on disk.

To rename a file, highlight it and select **Rename** from the **Shortcut** menu. You can also rename a file by selecting its name again after it has been highlighted.

To restore a file with a new name, highlight it and then select it again. Then type the new name.



Note that when you rename an instance, you are *only* renaming that file for the purposes of restoring it with this particular restore job. *Only the current restore job will assign that file the new name.* When you create a new restore job, you will see the file displayed with its original name. Similarly, the **Database** tab always displays files with the names they had when they were backed up.

Other Tips

Here are two additional tips for transferring files between operating systems and for setting up an autoloader for cleaning.

Moving Data between Operating Systems

- Clear the **Native data streams** check box in the **Advanced Options** window.


Suppose you want to transfer data (files and folders) from one operating system to another, such as from a NetWare platform to a Windows NT platform. To do so, you need to backup and restore the data in a generic format.

For further information on native data streams, see "Native data streams format," Chapter 7.

Different network software transmits data across the network to TapeWare in different formats. In particular, Windows NT and NetWare use different data stream formats. If you are going to share data from one LAN platform to another, the data should be stored on media in a common data format, *not* in the native data streams format.

To backup data in a generic format, first create a new backup job and select the data you want to transfer between operating systems. Then open the **Advanced Options** window and clear the **Native data streams** check box. When TapeWare backs up this data, it will convert it to a generic format before writing it to the media. After the job has completed, create a restore job, selecting the same files to restore. Make certain you have selected the proper instances of these files by selecting the proper instance date in the **Instances** window. You can then restore the files to a different operating system.

To transfer files between operating systems, clear the **Native data streams format** option on the backup job.



Setting up an Autoloader for Automatic Cleaning

- *Insert a cleaning cartridge into the autoloader and then change the status of the storage slot.*

Depending on the model and manufacturer, some autoloaders support automatic cleaning cycles. They will alert TapeWare when a cleaning cycle needs to be performed. If TapeWare knows that a particular storage slot on an autoloader magazine holds a cleaning cartridge, TapeWare will automatically run a cleaning cycle before running a backup job whenever a cleaning cycle is required.

For further information on setting up a cleaning cartridge on an autoloader, see "Status Tab," Chapter 12.

To set up an autoloader for automatic cleaning, insert the cleaning cartridge into the autoloader magazine. Then open the **Status** tab of the autoloader and change the status of the slot that holds the cleaning cartridge. Select the storage slot and then click the **Change Status** button. In the **Change Status** window, change the storage status to **Clean**. TapeWare will automatically use the cartridge in this slot when performing a cleaning cycle.

To manually clean an autoloader, highlight the device and select **Clean Device...** from the **Device** menu. Alternatively, open the autoloader's **Status** tab and select the device you want to clean. TapeWare will check to see if one of the slots holds a cleaning cartridge. If it does, the cleaning cycle will be performed in the background; if not, an error message is shown.

If you are using a device that is not an autoloader, you must manually clean the device at the manufacturers suggested intervals.

Permissions and Security Reference

This chapter provides a detailed summary of TapeWare's extensive security system. If it is your responsibility to manage the security of your TapeWare storage management database and you are working with sensitive data, this chapter can help you set up a complex security system that meets your particular security needs.

In This Chapter

- Overview
- Adding New Users and Groups
- Effective Permissions
- Permissions Reference

Overview

Permissions control what actions a user is allowed to perform within a given storage management zone. Users can be given extensive or limited permissions, allowing the TapeWare administrator to distribute backup duties to various users and workgroups. This allows for a flexible, non-centralized backup system while providing the highest degree of security for the network.

How your security is arranged depends on your unique security needs. Before setting up your security system, consider the following questions.

- *Are more than one storage management zone required?*

Setting up separate storage management zones can provide a high level of security. If your security needs require that access to some data be strictly limited, setting up a separate storage management database is often the simplest way to achieve this.

Data cannot be shared between zones without using advanced procedures. Media from one storage management database must be imported into a new database before the data on it can be read or used. When it is imported, TapeWare requires the media password. If you assigned the media a

password when it was created, the media cannot be imported without that password.

(On the other hand, if you do not assign the media a password, it can be easily imported into any storage management database. This actually leads to the data being less secure when there are more than one database than it would be with just one. Thus, if you are relying on multiple storage management databases for security purposes, be certain that all media created are assigned passwords.)

There may be, however, some limitations on the number of storage management databases you can set up. In particular, machines (workstations and file servers) can only be an object in one storage management database. Because tape drives and other backup devices are the peripherals of a machine, these devices can only be the member of one storage management zone as well. Similarly, volumes can only belong to one storage management zone. Files in one storage management zone cannot, without importing the media, be shared with database objects in other storage management zones.

Thus your ability to set up separate storage management zones is limited by the number of backup devices you have and their respective locations on separate machines. For example, to set up two storage management databases, you would require at least two separate workstations or file servers, each with at least one backup device.

- *Within a single storage management zone, must some users be prevented access to some data?*

Multiple workgroups may share a single tape drive or backup device, and thus are members of the same storage management zone. However, there may be reasons to allow these groups to work with only their own data. For example, an accounting workgroup may share a common tape drive with a personnel workgroup, although neither can be allowed access to the files and directories of the other group.

The security needs of these situations can be addressed by carefully assigning permissions, particularly to the machines, backup devices, media, volumes, and directories.

- *Should access to certain functions be limited?*

You may wish to distribute certain backup tasks to various users or workgroups. For example, each workgroup might be responsible for their own daily backup jobs and archive jobs. On the other hand, access to certain TapeWare features may need to be limited. Users might be able to *create* tapes, for example, but not *restore* files to disk or *delete* files on disk.

Alternatively, you may want users to *run* jobs you create, but not *create* their own jobs.

The security needs of these situations can be addressed by carefully assigning users select permissions to various objects in the database. For example, you might assign permission to write files to tapes, but not to volumes, thus preventing restore jobs from running.

Before Proceeding Further

The most powerful user in any storage management database is the TapeWare administrator. Because TapeWare administrators are granted supervisor rights to the system container, they have unlimited access to all of the objects in the storage management database. Any user who logs on as the TapeWare administrator will have complete access to all of the files and machines on the storage management database.

Your first security step should be to *change the TapeWare administrator's password*. Change your password on the **Security** tab. Select your User object and select **Change Password...** from either the **Security** menu or the **Shortcut** menu. Continue to change your password regularly and to use unique, carefully chosen passwords.

Before continuing with the rest of this chapter, be certain that you have changed your password.

Note the only difference between the TapeWare administrator (the **Admin**) and other users is that the TapeWare administrator has **Supervisor** rights to the root object in the TapeWare hierarchy (that is, to the **System container**). You may create additional TapeWare administrators if you want, as well as rename the **Admin**. However, do not delete the **Admin** unless you have assigned another user **Supervisor** or **Access** permission to the **System container**.

Warning Do NOT delete the TapeWare administrator without creating another user with **Access** permission to the **System container**.

Adding New Users and Groups

Generally, the first step to arranging the security system is to set up users and groups. You create new users and groups on the **Security** tab of the main TapeWare window. Use either the **Shortcut** menu or the **Security** menu to create new users and groups.

New User/Group Folders

Each time you add a new user or group object to the **Security** tab, TapeWare automatically creates a new User/Group folder in the **Home** folder with the same name as the new user or group. For example, if you create a new user named “Galen Clark,” TapeWare creates a new User/Group folder named “Galen Clark Folder.”

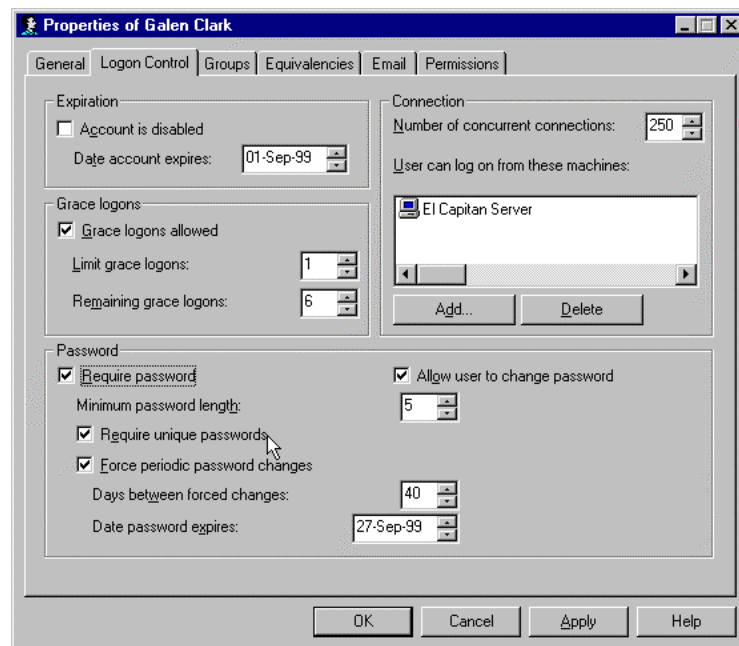
The user or group is automatically assigned six permissions to their User/Group folder: **Access**, **Create**, **Modify**, **Delete**, **Write**, and **Read**. You can modify these permissions at any time, including from the **Permissions** tab of the new user or group property sheet.

Setting Up Users

As you create new users, TapeWare automatically opens the property sheet of the new user. Use the tabs on the property sheet to control the user’s password, account activity, group membership, equivalencies, and permissions.

Logon Control Tab, User Object

The **Logon Control** tab controls whether passwords are required, whether and when the password must be changed, whether an account has expired, and the number of connections a user can have to the network.



The Logon
Control tab

Expiration: A user account can expire on a given date. When the account expires, TapeWare disables the account and checks the **Account is Disabled** box. This user will be unable to log on until the **Account is Disabled** box is cleared.

You can manually disable an account by checking the **Account is Disabled** box.

To make a disabled account active again, clear the **Account is Disabled** box and change the **Expiration Date**.

Connection: These parameters control from where a user can log on to TapeWare and how many connections a user can have simultaneously.

You can limit the number of a user's concurrent connections by changing the parameter in the **Number of concurrent connections** box. This parameter controls how many different logons a user may have simultaneously from different workstations or file servers. For example, if the **Number of concurrent connections** is set to 5, this user will be allowed to log on to TapeWare from up to five separate workstations or file servers at once.

Similarly, you can control from where a user can log on to TapeWare. The user will only be allowed to log on to TapeWare from the machines listed in the **User can log on from these machines** list box. To add machines, click **Add...** and select the appropriate machine from the **Browse** window. Note that if no machine is listed, a user can log on from any machine.

Password: When **Require password** is checked, TapeWare requires the user to have a password. The minimum length of the password is determined by the **Minimum password length** parameter. TapeWare will check to see if the password is unique when **Require unique passwords** is checked.

You can force the user to change their password regularly by checking the **Force periodic password changes** box and the **Allow user to change password** box.

Note that when you clear the **Require password** check box, if the user still has a password, TapeWare will continue to require that the user input their password.

Tip The TapeWare administrator can change a user's password without knowing the user's current password. When the TapeWare administrator selects **Change Password...** from the **Security** or **Shortcut** menu, TapeWare does not require that the old password be entered before changing the password. This is useful when the user has forgotten his or her password.

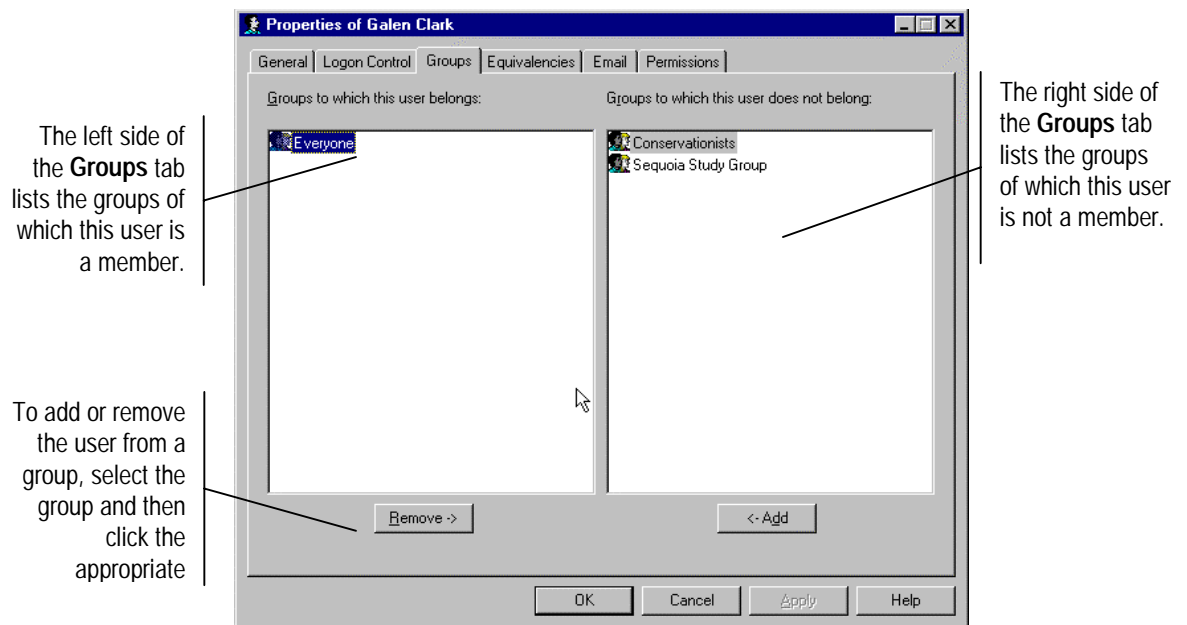
Grace Logons: If the **Force periodic password changes** box is checked, TapeWare will prompt the user to change their password as the user logs on. **Grace Logons** refer to the number of times a user can log on to TapeWare with their old password when that password has expired. For example, if the **Allow**

Grace Logons box is checked and the number of grace logons is set to 2, the user will be allowed on log on two times using their old password, even though that password has expired. The third attempt to log on will be denied.

Note that grace logons do not function when passwords are not required, that is, when the **Require password** check box is cleared.

Groups Tab, User Object

Use this tab to add or remove a user from a group. To add a user to a new group, select the group on the right side of the window and click the **Add** button; the group will move to the left side of the window. Similarly, to remove a user from a group, select the group on the left side of the screen and then click the **Remove** button.



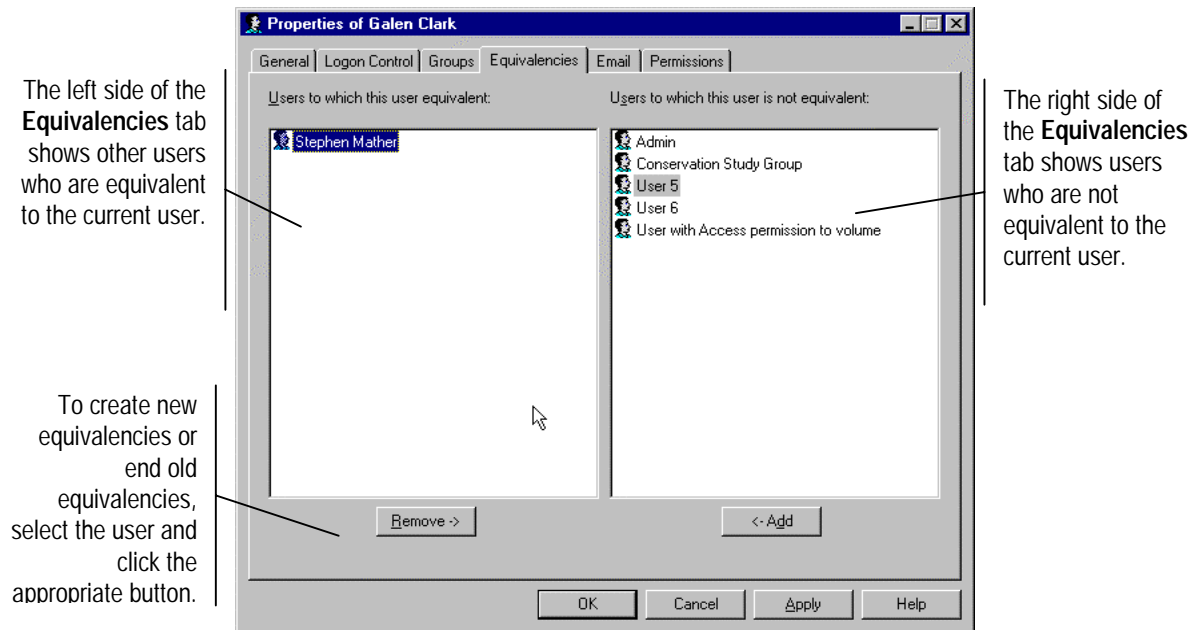
Everyone group: When a new user is created, they are automatically added to the **Everyone group**. Members of this group have **Modify**, **Delete**, **Create**, **Write**, and **Read** permissions to the **Everyone Folder**. You can modify these permissions at any time, including from the **Permissions** tab of the new user's property sheet.

Equivalencies

One quick way to assign a user permissions is to make the current user equivalent to another user. This can be very useful for managing complex

TapeWare installations with multiple users and varying security or for making temporary changes to a user's permissions.

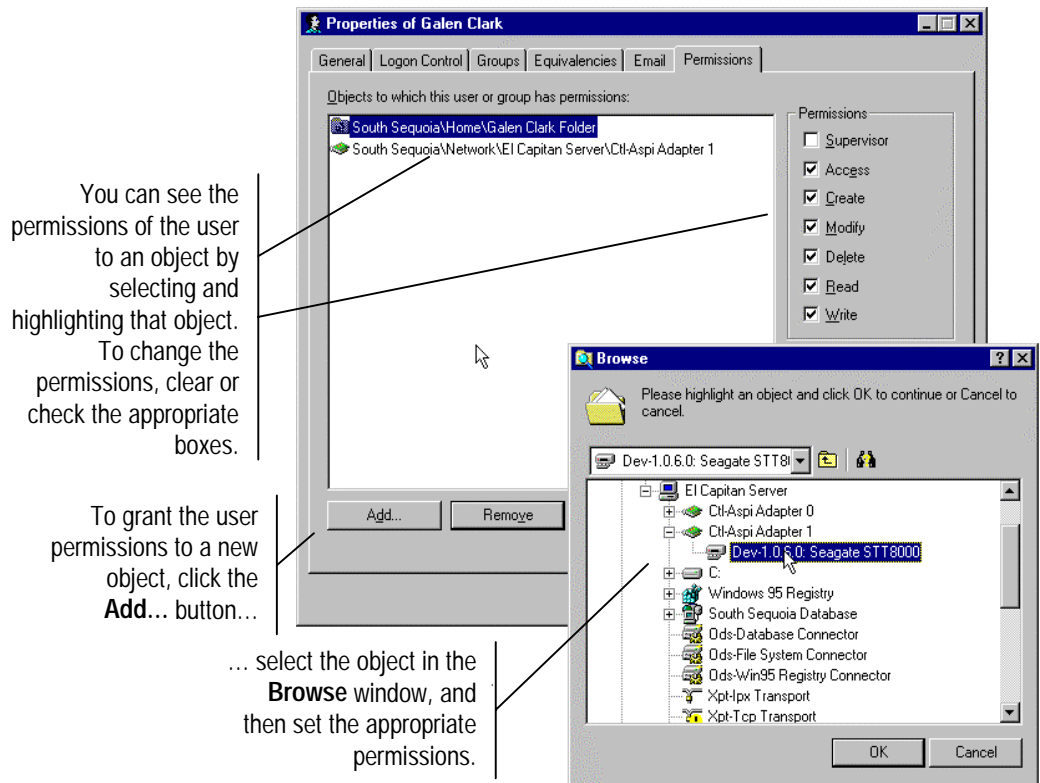
Use this tab to make the current user equivalent to another user. To make the current user equivalent to another user, select the other user on the right side of the window and click the **Add** button; the user will move to the left side of the window. Similarly, to end an equivalency, select the other user on the right side of the screen and then click the **Remove** button.



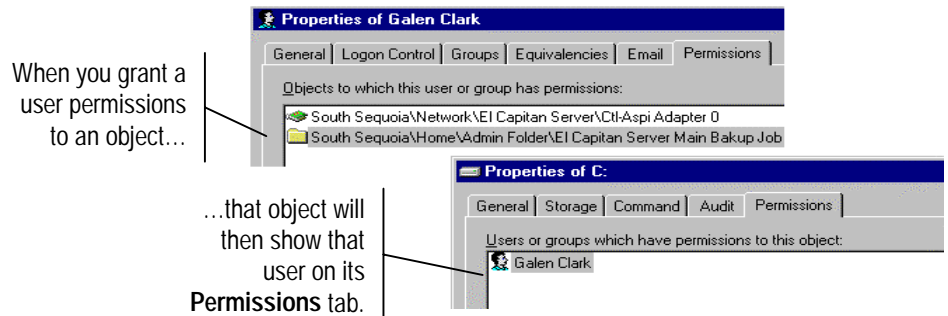
Note that equivalencies only work in one direction; they are not reciprocal. The current user's effective permissions (the effective permissions of the user whose property sheet is open) will be calculated using the direct and inherited permissions of the user they are made equivalent to. For example, if a user named Galen Clark is made equivalent to a user named Stephen Mather, Clark's effective permissions are calculated using *both* Mather's *and* Clark's direct permissions. However, Mather's effective permissions remain unchanged.

Permissions Tab, User Object

Use this tab to grant users permissions to objects in the storage management database. The **Permissions** check boxes show the permissions of whatever object is selected in the **Objects to which this user or group has permissions** list. Select another object to see the user's permissions to that object.



Note that permissions can be granted from either the property sheet of the database object or the property sheet of the user. Either way, the permissions appear on the appropriate corresponding object's **Permissions** tab. For example, if Galen Clark is granted permissions to the **C:** volume from the **Permissions** tab on his property sheet, the **Permissions** tab on the property sheet of the **C:** volume will list Clark as a user who has permissions. Alternatively, if Clark is granted permissions from the property sheet of the **C:** volume, the appropriate permissions will appear on Clark's **Permissions** tab.



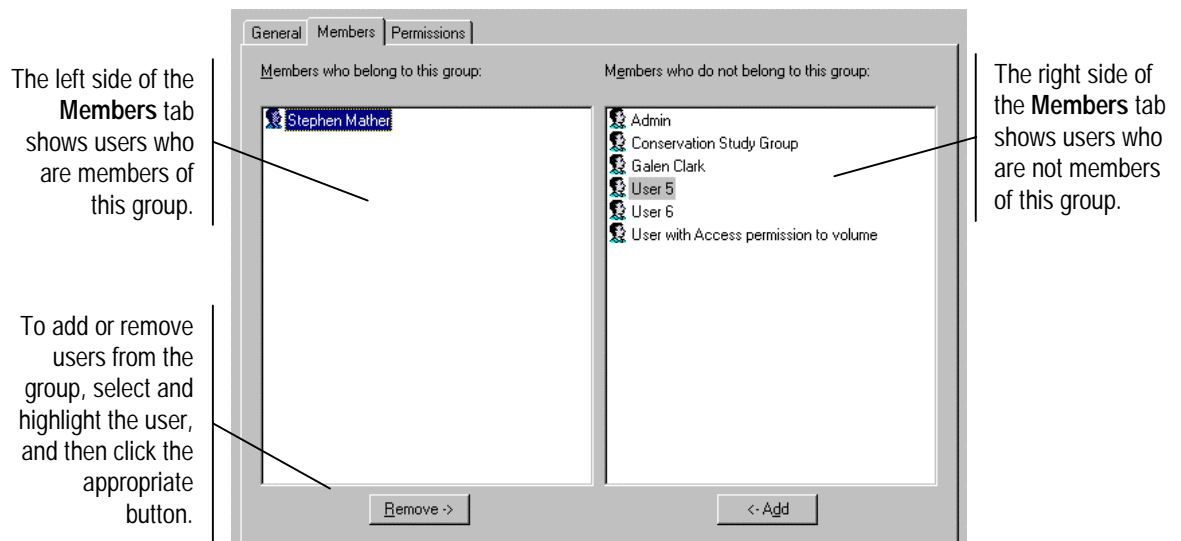
Note additionally that a user has direct permissions only to those objects listed on that user's **Permissions** tab. Any and all other effective permissions to other objects are calculated through inherited permissions, through equivalencies, or through groups.

Setting Up Groups

As you create new groups, TapeWare automatically opens the property sheet of the new group. Use the tabs on the property sheet to assign members to the group and assign permissions to the group.

Members Tab

Use the **Members** tab of the group's property sheet to add and remove users from the group. To add a user to the group, select the user on the right side of the window and click the **Add** button; the user will now appear on the left side of the window, under **Members who belong to this group**. To remove a user from a group, select the user on the left side of the window and click the **Remove** button; the user will move to the right side of the window, under **Members who do not belong to this group**.



Permissions Tab, Group Object

If a user is a member of a group, that user's effective permissions are determined using the direct permissions that group has to objects in the storage management database. Use this tab to assign the group permissions to objects. The **Permissions** check boxes show the permissions of whatever object is selected in the **Objects to which this user or group has permissions** list. Select another object to see the Group's permissions to that object.

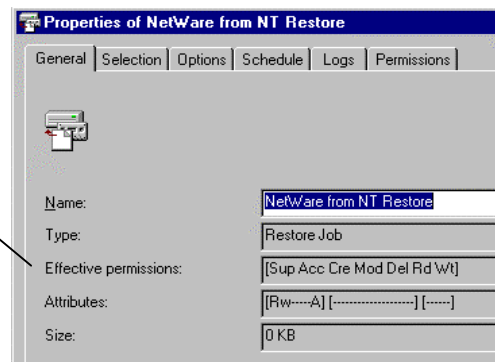
Note that the permissions granted from this tab, like all permissions, are reciprocal. Changes made on this tab appear on the tabs of the corresponding object. For example, if you grant a group permissions to a folder, the **Permissions** tab of that folder will list the group, along with the appropriate corresponding permissions.

Effective Permissions

TapeWare insures the security of the storage management database and LAN by calculating the **effective permissions** a user has to an object and using these permissions to determine what actions that user can perform.

The current user's effective permissions to an object are displayed on the **General** tab of the object's property sheet. The **Effective Permissions** box shows the current user's effective permissions to the object.

The current user's effective permissions to an object are shown on the **General** tab of that object.



Calculating Effective Permissions

A user's effective permissions are calculated using either the user's **direct permissions** or the user's **inherited permissions** (but *never both* the direct and the inherited permissions.)

A user has *direct permissions* to an object as a result of three situations: (1) the user is listed on the **Permissions** tab of the object (and reciprocally, the object is listed on the user's **Permissions** tab); (2) the user is equivalent to a user who has direct permissions to the object; and (3) the user is a member of a group that has direct permissions to the object. Note that these three ways of gaining permissions are not mutually exclusive: a user has direct permissions in only one of the ways, in two of the ways, or in all three of the ways.

A user has *inherited permissions* to an object only if *both* (1) the user does not have direct permissions to the object, *and* (2) the user has effective permissions

to the container that contains the current object. The user's effective permissions to the container object can be either direct or inherited permissions.

Effective Permissions Algorithm

TapeWare uses the following algorithm to determine effective permissions.

- Does the user have direct permissions to the object? If yes, these are used to calculate the effective permissions. TapeWare does not check to see if the user has inherited permissions.
- Does the user have effective permissions to the container that contains the current object (inherited permissions)? If yes, these permissions are used to calculate the effective permissions. If not, then the user does not have effective permissions to the object.

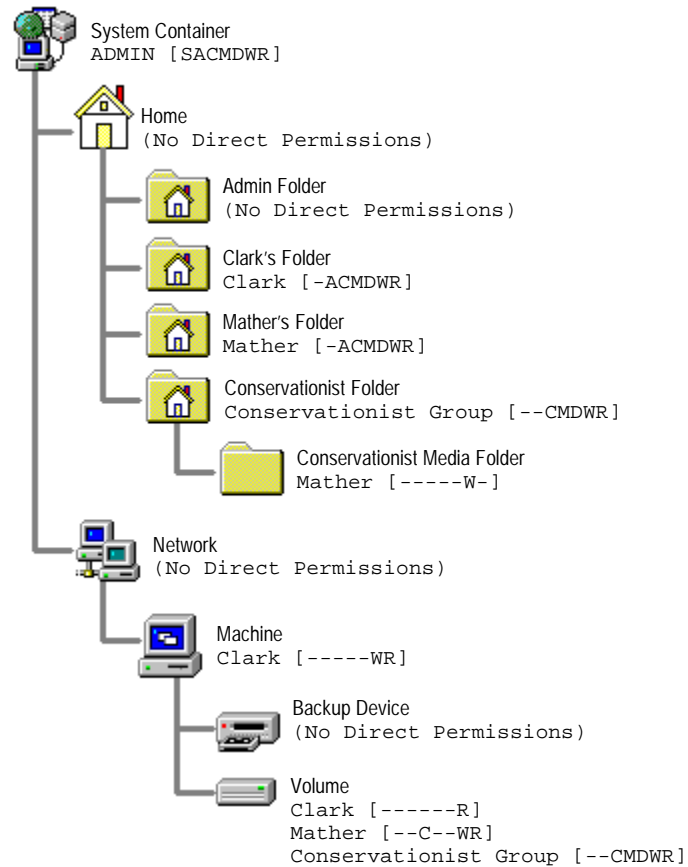
Permissions from Multiple Sources

Users can gain *direct* permissions to objects either as users, as a result of equivalencies, or as members of a group. When the direct permissions result from multiple sources, TapeWare uses all of the sources to determine the permissions.

Consider the following example: Galen Clark has direct permissions to **Read** and **Write** to a folder called **Conservationist Archive Jobs**; Clark is also a member of a group that has direct permissions to **Modify** to the folder. As a result, Clark's effective permissions are **Read, Write, and Modify**.

Examples of Effective Permissions

The following six examples illustrate how effective permissions are calculated. The diagram below illustrates these six examples.



- (1) The TapeWare administrator has direct permissions to the **System Container**, the object at the very top of the storage management database hierarchy. These determine his or her effective permissions to this object. Because it is a container, the objects below it in the storage management database all have inherited permissions because the object directly above them has effective permissions. So, for example, the TapeWare administrator has effective permissions to the **Home Folder** because it inherits its permissions from the object that contains it, the **System Container**. Thus, the TapeWare administrator has effective permissions to all of the objects in the storage management database.

- (2) A user named Stephen Mather has direct permissions to his User/Group folder, named **Mather's Folder**. As a result, by inherited permission, Mather has effective permissions to the objects stored in this folder, including any jobs, media, or job folders stored in this folder. Mather does not, however, have effective permissions to the **Home Folder** or to the **System Container**—these objects are *above* his User/Group folder and thus do not inherit permissions.
- (3) A user named Galen Clark has direct permissions to a **Machine**, in this case a file server with an attached tape drive and several associated disk drives. The direct permissions to the file server mean that Clark also has effective permissions (by inheritance) to the tape drive. So for example, Clark might be given read and write permissions to the file server, and thus to the tape drive.

However, Clark is prevented from having permissions to the volumes on the file server. He is listed on the **Permissions** tab of the volume and these direct permissions are used to deny him access to the volume. In this example, he is granted **Read** permission by checking that box, but denied **Write** permissions by clearing the appropriate box.

Thus even though Clark has effective permissions to the container that contains the volume, Clark's effective permissions to the volume are determined *only* by his direct permissions to the volume. Because Clark has direct permissions, TapeWare does not check to see if Clark has inherited permissions.

- (4) The following example is more complex, but illustrates an important concept: that TapeWare does not check for inherited permissions when there are direct permissions.

A user named Stephen Mather is a member of the **Conservationist** group, which has five direct permissions to the **Conservationist Folder**: Create, Modify, Delete, Write, and Read permissions. Mather also has direct permissions to the **Conservationist Media Folder**, but only Write permission.

Mather has five effective permissions to objects contained in the **Conservationist Folder**, but not to the **Conservationist Media Folder**, where he has only one (Write permission). TapeWare does not look to see if Mather has effective permissions to the container that contains the **Conservationist Media Folder** because Mather has direct permissions to that object. Thus even though other members of the **Conservationist** group have effective permissions to the **Conservationist Media Folder** through inherited permissions, Mather will not. Mather will have only Write permissions to this folder.

- (5) The following example shows how equivalencies and group membership work together to determine effective permissions.

Suppose that Mather is a member of the **Conservationist** group *and* that he is made equivalent to Clark. What permissions will Mather have?

Mather has permissions to all of the User/Group folders, except the **Admin Folder**. For example, he has permissions to **Clark's Folder** because he is equivalent to Clark. (Note that this equivalency does not give Clark permission to **Mather's Folder**.) Mather also has the same permissions to the **Machine** and **Tape Drive** that Clark does.

However, Mather's permissions to the **Volume** are different from Clark's. Mather has direct permission to the **Volume** in three ways: as a user, as a member of the **Conservationist** group, and as a result of his equivalency to Clark. When TapeWare calculates his effective permissions, it uses these direct permissions from all three sources. In this case, Mather will have five permissions (Create, Modify, Delete, Write, and Read).

Note that it does not matter that Mather's own direct permissions as a user do not include Create and Modify permissions. TapeWare uses all three sources to determine Mather's effective permissions to the volume. In this case, Mather's membership in the **Conservationist** group grants him Create and Modify permissions.

- (6) Given the above example, suppose we wanted to deny Mather *all* permissions to the **Volume**. How could this be accomplished?

To deny Mather all permissions to the **Volume**, three things must happen: Mather's equivalency to Clark must end; Mather's membership in the Conservationist group must end; and Mather's direct permissions must be changed so that Mather is listed on the **Permissions** tab of the **Volume** but no permission boxes are checked.

Note that listing Mather on the **Permissions** tab and clearing the permissions check boxes is not enough to deny Mather permissions to the tab. Mather must no longer be equivalent to Clark and Mather must no longer be a member of the Conservationist group.

Checking Effective Permissions

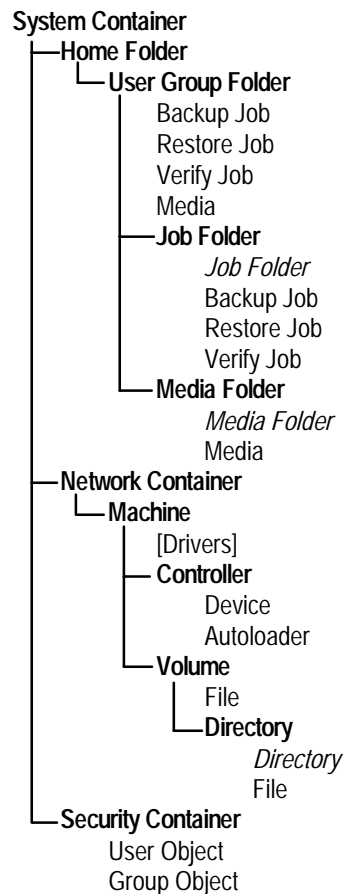
On complex installations with multiple users and groups and varying levels of security, the effective permissions a particular user has can be difficult to intuit.

There is a simple way to check what effective permissions a particular user has: by logging on as the user.

First, if necessary, you can temporarily disable the password by clearing the **Require password** check box on the user's **Logon Control** tab. Then, log on as the user. Browse the various **General** tabs of the objects in the storage management database and check to see if the effective permissions displayed match your intended security measures.

Permissions Reference

There are seven permissions: **Read, Write, Delete, Modify, Create, Access,** and **Supervisor**. These permissions effect different objects in the TapeWare storage management database differently. Even though a particular permission may not apply directly to that object, objects below it in the storage management database hierarchy can still inherit permissions from that object.



Bold = Container
Italic = Container stored in
 container of same type.

Read Permission

Affected Objects: *Media, Controller, Device, Autoloader, Volume, Directory, File*, plus storage management database.

Description: Controls whether a user can read from a given storage management database object.

In the case of physical peripherals that perform read functions, such as Controllers, Devices, Autoloaders, and Volumes, **Read** permission to the peripheral is required in order for TapeWare to instruct the peripheral to read files or directories.

In case of storage management database objects that hold data, such as Media, Volumes, Directories, and Files, **Read** permission is required to read the data these objects contain.

Affected: *Commands Copy, Run (job type), Rewind, Start, Eject Media, Eject Magazine, Retension, Restore Database, Clean Device, Identify Media, Import Media, Restore Database*

This permission enables **Copy** (but not **Paste**), allowing the user to copy objects in the storage management database.

Read: permission is also required to run jobs. Backup jobs require **Read** permission to the appropriate Volumes, Directories, and Files; restore jobs require **Read** permission to the appropriate Devices, Autoloaders, and Media; verify jobs require **Read** permission to all of these objects.

Many commands that perform utility functions, such as **Clean Device** or **Eject Media** on a device require **Read** permission. Device commands that also read media in backup devices require this permission.

Write Permission

Affected Objects: *Media, Device, Autoloader, Volume, Directory, File*, plus database.

Description: Controls whether a user can write to a given storage management database object.

In the case of physical peripherals that perform write functions, such as Controllers, Devices, Autoloaders, and Volumes, **Write** permission to the peripheral is required in order for TapeWare to instruct the peripheral to write files or directories.

In the case of storage management database objects that hold data, such as Media, Volumes, Directories and files, **Write** permission is required to write data to these objects.

Affected Commands: *Run (job type)*

Write: permission is also required to run jobs. Backup jobs require **Write** permission to the appropriate Devices, Autoloaders, and Media; restore jobs require **Write** permission to the appropriate Volumes, Directories, and Files.

Delete Permission

Affected Objects: *Media, Device, Autoloader, Volume, Directory, File*, plus database.

Description: Controls whether a user can delete storage management database objects or perform delete functions.

In the case of physical peripherals that perform delete functions *including overwrite functions*, such as Controllers, Devices, Autoloaders, and Volumes, **Delete** permission to the peripheral is required in order for TapeWare to instruct the peripheral to delete or overwrite files or directories.

In case of database objects that hold data, such as Media, Volumes, Directories, and Files, **Delete** permission is required to delete or overwrite the data these objects contain.

Affected Commands: *Delete, Run (job type)*

This permission enables **Delete**, allowing the user to delete objects in the storage management database.

Delete: permission is also required for some types of jobs. Backup jobs require **Delete** permission to the appropriate Devices, Autoloaders, and Media whenever files are overwritten or media formatted; restore jobs require **Delete** permission to the appropriate Volumes, Directories, and Files whenever the files are overwritten.

Modify Permission

Affected Objects: All storage management database objects.

Affected Property Sheet Tabs: *General, Selection, Options, Schedule, Logs, Storage*, plus machine diagnostic tabs.

Description: Controls whether a user can modify the specified tabs on an object's property sheet. Controls whether a user can change the name of an object.

Controls whether a user can move an object to a new location in the storage management database.

For any object, **Modify** permission allows the user to change the object's **General** tab. (This effects only the name of the object.)

For Backup Jobs, Restore Jobs, and Verify Jobs, **Modify** permission *to the job* allows the user to change the job's **Selection**, **Options**, **Schedule**, and **Logs** tabs. Note that **Read** permission to the Volume is required in order to select the Volume's files and directories on the **Selection** tab.

For Machines, Controllers, and Volumes, **Modify** permission allows the user to modify the diagnostic tabs, such as **Communication Test**, **Ping Test**, etc.

Affected Commands: *Move...*, *Rename*

This permission enables **Move...** and **Rename**, allowing users to move objects in the storage management database and rename them.

Create Permission

Affected Objects: *Home Folder*, *User/Group Folder*, *Job Folder*, *Media Folder*, *Restore Job*.

Description: Controls whether a user can create new objects within a container object.

For the Home Folder, **Create** permission is required in order to create new User/Group Folders.

For User/Group Folders, **Create** permission is required in order to create new Job Folders, Backup Jobs, Restore Jobs, Verify Jobs, Media Folders, and Media.

For Job Folders, **Create** permission is required in order to create new Job Folders, Backup Jobs, Restore Jobs, and Verify Jobs.

For Media Folders, **Create** permission is required in order to create new Media Folders and Media.

For Restore Jobs, jobs that restore files in new locations or with new names require **Create** permission to the appropriate machines, volumes, and directories.

Affected Commands: *New...Job*, *New...Folder*, *New Object...*, *Paste*.

This permission enables the **New...Job** and **New...Folder** commands, for each type of job and folder. This permission also enables the **New Object...** command on the **File** menu.

This permission enables **Paste**, allowing the user to paste objects in the storage management database.

Access Permission

Affected Objects: All storage management database objects, except *Security Container*, *User Object*, and *Group Object*.

Effected Property Sheet Tabs: *Permissions*.

Description: Controls whether a user can see and modify the **Permissions** tab of an object.

For any object, **Access** permission allows the user to change the permissions to the object. To add a new user to the **Permissions** tab, **Modify** permission is also required. A user listed on the **Permissions** tab can be deleted with **Access** permission alone.

Note that **Access** permissions does not allow the user to change the **Permissions** tab on either the Security Container, a User Object, or a Group Object.

Supervisor Permission

Affected Objects: All storage management database objects.

Effected Property Sheet Tabs: *Logon Control*, *Equivalencies*, *Groups*, *Members*, *Permissions*

Description: This permission gives the user unlimited permissions to the object and all objects below it in the storage management database. Additionally, only a user with **Supervisor** permission to the Security Container can create new users and groups.

When a user has **Supervisor** permission to an object, the user is automatically granted all seven permissions to the object. Furthermore, the user cannot be denied any permission to any object below it in the hierarchy, even by assigning that user direct permission. Thus, as a result, *a user with supervisor permission to an object will have all permissions to the object and every object below it in the storage management database.*

Supervisor permission to the Security Container is required to create new users and groups. Additionally, the **Logon Control**, **Equivalencies**, **Groups**, and **Members** tabs are only available to users with **Supervisor** permission to the Security Container.

Normally, the TapeWare administrator is given **Supervisor** permission to the System Container.

Affected Commands: *New User, New Group*

This permission enables the **New User** and **New Group** commands, allowing the user to create new users and groups.

Objects and Properties Reference

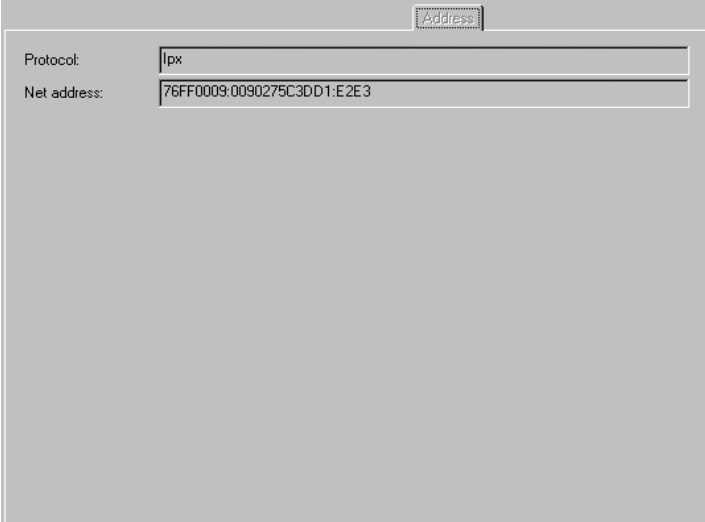
This chapter provides reference details for every object and property in the TapeWare database. It is organized alphabetically according to the name of each property sheet tab. The property sheet tab name appears at the beginning of each entry and in the page header. The **Applicable Objects** heading lists the objects this property sheet tab applies to. Various fields and parameters on each property sheet tab are indicated by bold headings, with a short description following each. List box choices and field parameters are indicated by bold in-line headings.

In addition to the property sheet tabs, this chapter provides reference information for the **Instances** window, the **Preferences** window, and the **Selection Filters** window.

Address Tab

Applicable Objects: *Machine*

Shows the machine object's network address.

A screenshot of a network configuration window titled "Address". The window has a light gray background and a thin border. At the top, the title "Address" is in a small box. Below the title, there are two labels: "Protocol:" and "Net address:". To the right of "Protocol:" is a text box containing "Ipx". To the right of "Net address:" is a text box containing "76FF0009:0090275C3DD1:E2E3". The rest of the window is empty.

Protocol:	Ipx
Net address:	76FF0009:0090275C3DD1:E2E3

Address Tab

Protocol

The protocol is either IPX or TCP/IP, depending on your network operating system.


Net Address

Shows the machine address used by the network operating system.

Audit Tab

Applicable Objects: *File, Directory, Volume*

Shows the audit log for the object.



Audit Tab

The audit tab appears on the property sheet of an object that has been *audit enabled*. To enable an objects audit log, open the **Storage** tab of the object and change the **Audit object actions** setting to **Audit enable**.

When audit enabled, TapeWare enters into that object's log a record every action performed on that object. For example, there will be an entry each time the object is backed up or restored. Additionally, the audit log will show the media on which instances of the object are stored.

You can print or save the audit log using the text editor specified on the **Preferences** tab.

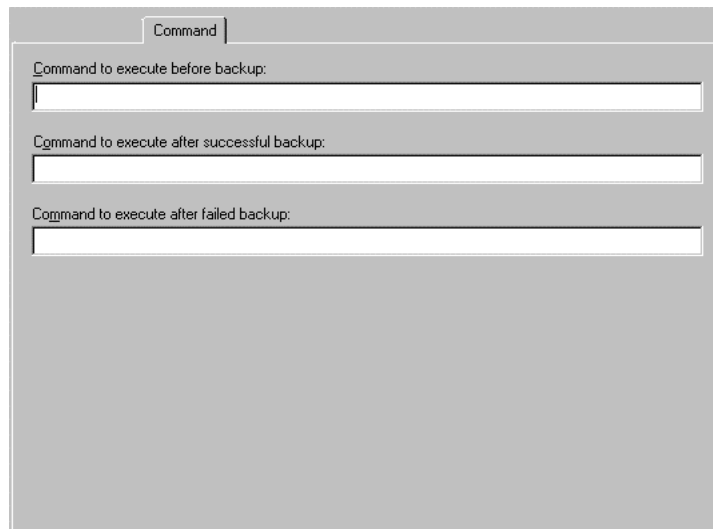
Command Tab

Applicable Objects: *Volume, Directory, File*

Use this tab to send and execute operating system commands before and after backup jobs. You might use this command, for example, to close a large database file before backing it up and then to open it again after it was successfully backed up.

The default path is the same as the path of the current object (e.g., the volume, directory, or file whose property sheet is open). You can specify another path if necessary.

Specified commands must be executable by the operating system. This includes .bat and .ncf files. Commands are operating system specific.

The image shows a screenshot of a software window with a tab labeled "Command". Inside the window, there are three text input fields. The first field is labeled "Command to execute before backup:". The second field is labeled "Command to execute after successful backup:". The third field is labeled "Command to execute after failed backup:". Each field has a small cursor icon at the beginning. The background of the window is a light gray.

Command Tab

Command to execute before backup

This command is sent on the path of the current object before that object is opened for backup.

Command to execute after successful backup

This command is sent on the path of the current object after that object has been successfully backed up and closed.

Command to execute after failed backup

This command is sent on the path of the current object after that object has been closed after an unsuccessful attempt to backup that object.

Communication Test Tab

Applicable Objects: *Machine*

This tab can be used to test communication layer of a network. Can be used to evaluate the ability to transfer data under optimum circumstances.

The screenshot shows a window titled 'Comm Test'. It features a 'Transfer type:' dropdown menu with 'Backup' selected. To the right of the dropdown is a checked checkbox labeled 'With data checking' and a 'Start' button. Below these is a 'Communication status' section containing six rows, each with a label and an input field: 'Packet No.', 'Transfer rate (KBytes/sec)', 'Transfer rate (MBytes/min)', 'No. of ACKs sent', 'No. of NACKs sent', and 'No. of timeouts'.

Communication Test Tab

Transfer type

Specifies the communication test to be performed.

Backup: Simulates data transfer during a backup job. Local machine sends large packet; remote machine sends small reply acknowledging receipt of data.

Restore: Simulates data transfer during a restore job. Remote machine sends large packet; local machine sends small reply acknowledging receipt of data.

Large Packet: Large-sized packet of data sent between remote and local machine.

Medium Packet: Medium-sized packet of data sent between remote and local machine.

Small Packet: Small-sized packet of data sent between remote and local machine.

Communication Status

Returns results from the performed test.

(Note the relative difference packet size makes in transfer rate. You can use this information to optimize backup job performance.)

With data checking

Known byte pattern is sent between machines. When this option is checked, this pattern is checked by machine receiving it.

For more information on...	See...
-----------------------------------	---------------

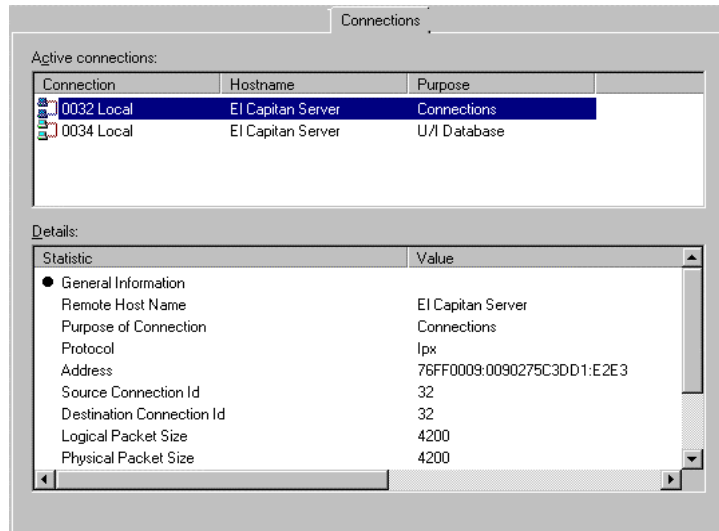
Optimizing backup jobs by increasing data transfer rate	“Strategies for faster jobs,” Chapter 10
---	--

Connections Tab

Applicable Objects: *Machine*

This descriptive tab shows the active connections for this machine. It is for information purposes only.

Connections are established for varying purposes and are automatically opened and closed as necessary.



Connections Tab

Active Connections

Lists the currently active connections for this machine.

Details

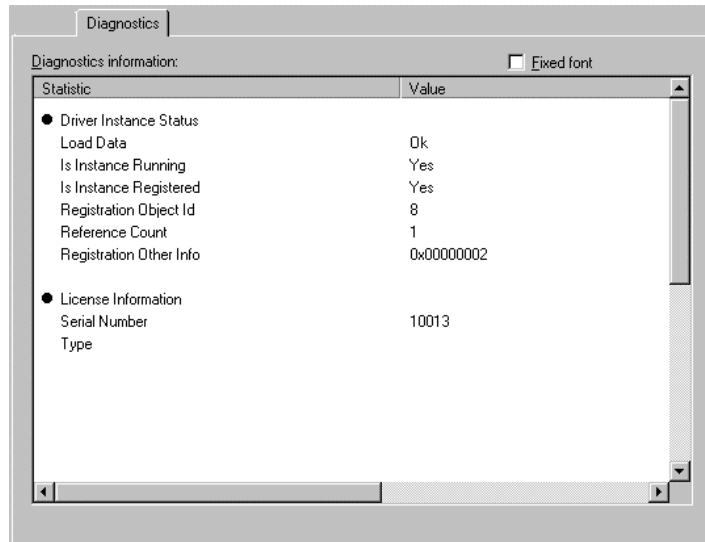
Shows detailed information about the selected active connection.

Diagnostics Tab

Applicable Objects: *Machine, Drivers*

Provides detailed information about the current machine or driver object. For information purposes only.

The **Diagnostics** tab is available for all active machines and drivers, including controllers, logical tape formats, and services.



Diagnostics Tab

Drivers Tab

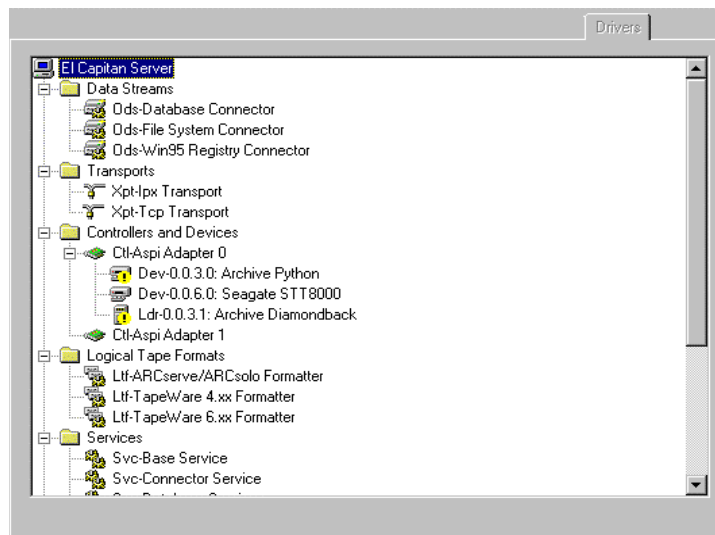
Applicable Objects: *Machine*

Shows the available drivers for this machine. For information purposes only.

Provides user with single view of all drivers associated with this machine. (Drivers also appear on **Database** tab in various locations.)

Drivers are used by TapeWare for differing purposes. The name of each folder on this tab indicates the purpose of the drivers in that folder.

Drivers which are marked with a yellow exclamation point are currently unavailable for some reason. For example, a controller driver may be marked with this icon when the device it is associated with is off-line. To restart the driver, you must correct the problem and then exit and restart TapeWare. When TapeWare restarts, it will reinitialize these devices.

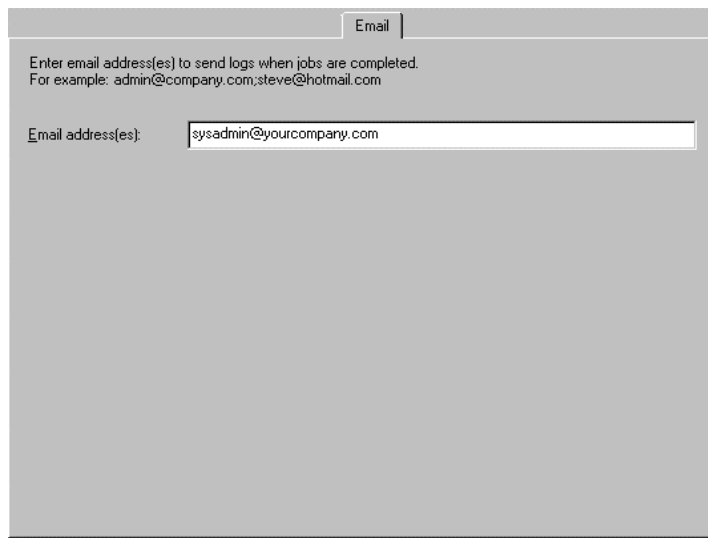


Drivers Tab

Email Tab

Applicable Objects: *User Object*

Shows the email address to which logs of jobs will be emailed when that user is the owner of a job that runs. This tab only appears when the optional Email package has been installed.



Email Tab

When Email has been installed and configured, TapeWare will attempt to email the log of every job after it has run. The log is sent to the job's **Owner**.

You can enter multiple email addresses by separating each address with a semicolon (no additional spaces).

The address entered here need not be the address of the user. For example, you can enter the TapeWare administrator's address in this field for any or all users.

For more information on...

See...

The owner of a job

"How Forcing Jobs to Run Effects Permissions," Chapter 8

Installing and configuring email

"Configuring Email," Appendix II

Using Email to send job logs

"Email Support for Job Logs," Chapter 8

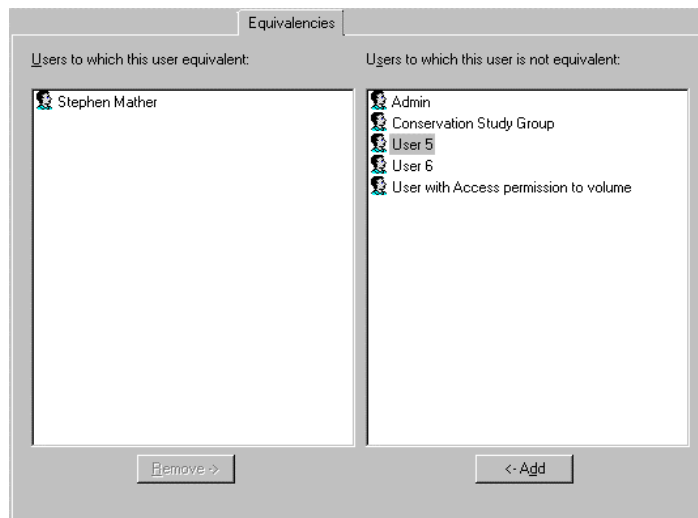
Equivalencies Tab

Applicable Objects: *User Object*

This tab is used to assign an individual user effective permissions equivalent to the effective permissions of another user.

Note that equivalencies only work in one direction; they are not reciprocal. The current user's effective permissions (the effective permissions of the user whose property sheet is open) will be calculated using the direct and inherited permissions of the user they are made equivalent to. For example, if a user named Galen Clark is made equivalent to a user named Stephen Mather, Clark's effective permissions are calculated using *both* Mather's *and* Clark's direct permissions. However, Mather's effective permissions remained unchanged.

Additionally, note that users can gain *direct* permissions to objects either as users, as a result of equivalencies, or as members of a group. When the direct permissions result from multiple sources, TapeWare uses all of the sources to determine the permissions.



Equivalencies Tab

Users to which this user is equivalent

Lists those users to whom the current user (whose property sheet is open) is equivalent. To remove users from this field, select them and then click **Remove**. To add users to this field, select them on the right side of the window and then click **Add**.

Users to which this user is not equivalent

Lists those users to whom the current user (whose property sheet is open) is *not* equivalent.

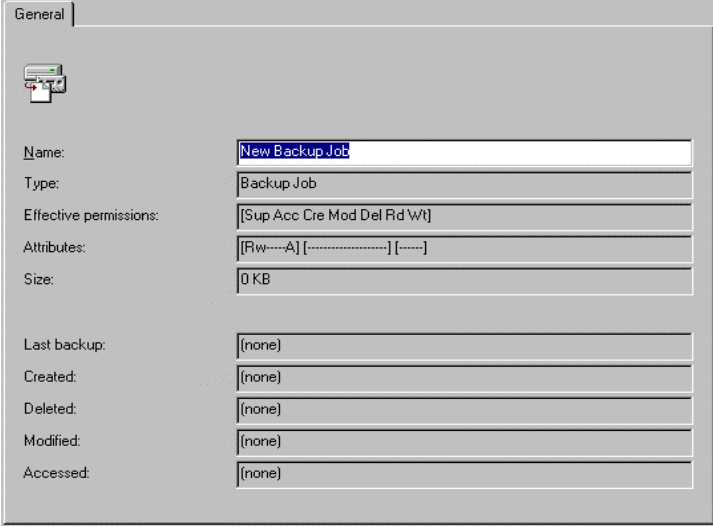
For more information on...	See...
Calculating effective permissions	“Effective Permissions,” Chapter 11
Examples of equivalencies	“Examples of Effective Permissions,” Chapter 11

General Tab

Applicable Objects: *All objects*

This tab shows information and object attributes stored in the TapeWare database for the current object.

The data on this tab is taken from the TapeWare database. For files, directories, and volumes, this data is regularly updated. Each time TapeWare opens a directory or volume, it updates its database with any new information about these files and directories.



The screenshot shows a window titled 'General' with a tab icon. It contains the following fields:

Name:	New Backup Job
Type:	Backup Job
Effective permissions:	[Sup Acc Cre Mod Del Rd Wt]
Attributes:	[RW-----A] [-----] [-----]
Size:	0 KB
Last backup:	(none)
Created:	(none)
Deleted:	(none)
Modified:	(none)
Accessed:	(none)

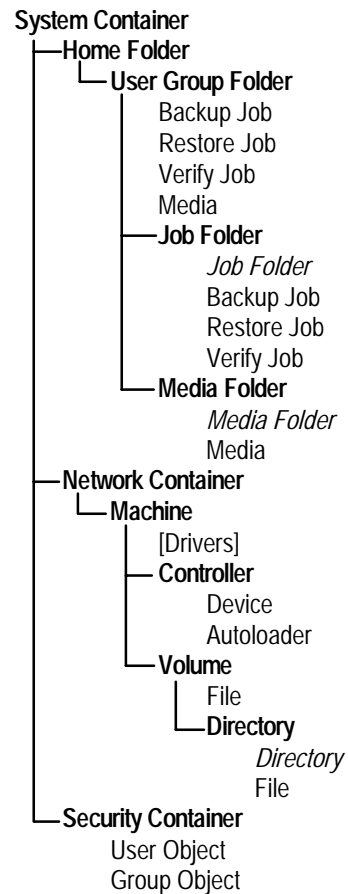
General Tab

Name

Shows the name of the current object. To change the name, select it and type in a new name. An object name can be up to 256 characters long.

Type

Shows the type or category of object. There are 21 object types in the TapeWare database, plus drivers. The object names and their possible relationships to each other are illustrated below.



Bold = Container
Italic = Container stored in
 container of same type.

Effective Permissions

Shows the effective permissions the current user has to this object. Note that these are the *effective* permissions, not the *direct* permissions. There are seven possible permissions, abbreviated as follows:

Sup	Supervisor
Acc	Access
Cre	Create
Mod	Modify
Del	Delete
Rd	Read
Wt	Write

Attributes

Shows operating system attribute information about file or directory. The attribute information is taken from the operating system and is updated each time the directory that contains the file is opened. You can use this information to sort files with filters.

The attribute abbreviations are as follows:

Ro	Read only
Rw	Read and write
H	Hidden
Sy	System
X	Execute only
D	Directory
A	Archive
Sh	Share
Tm	Temporary
T	Transaction
Ra	Read audit
Wa	Write audit
P	Immediate purge
Ri	Rename inhibit
Di	Delete inhibit
Ci	Copy inhibit
Dm	Migrate inhibit
Ds	Sub-allocation inhibit
Ic	Immediate compression
Dc	Don't compress
Co	Compressed
Cc	Can't compress
Mg	Migrated

Size

Shows operating system information about file or directory size. Information about the size of the file or directory is taken from the operating system and is updated each time the directory that contains the file is opened. You can use this information to sort files with filters.

Created

Shows operating system information about date file or directory created. The create date is taken from the operating system and is updated each time the directory that contains the file is opened. You can use this information to sort files with filters.

Deleted

When a file that was previously backed up is deleted, TapeWare assigns it a delete date. When TapeWare opens a directory, it compares the files it finds with information about instances of files in its database. When an instance of the file is found in the database, but not in the directory, TapeWare assigns that instance a delete date.

Modified

Shows operating system information about date file or directory last modified. The modify date is taken from the operating system and is updated each time the directory that contains the file is opened. You can use this information to sort files with filters.

Accessed

Shows operating system information about the date the file or directory last accessed. The access date is taken from the operating system and is updated each time the directory that contains the file is opened. You can use this information to sort files with filters.

For more information on...**See...**

Calculating effective permissions

“Effective Permissions,” Chapter 11

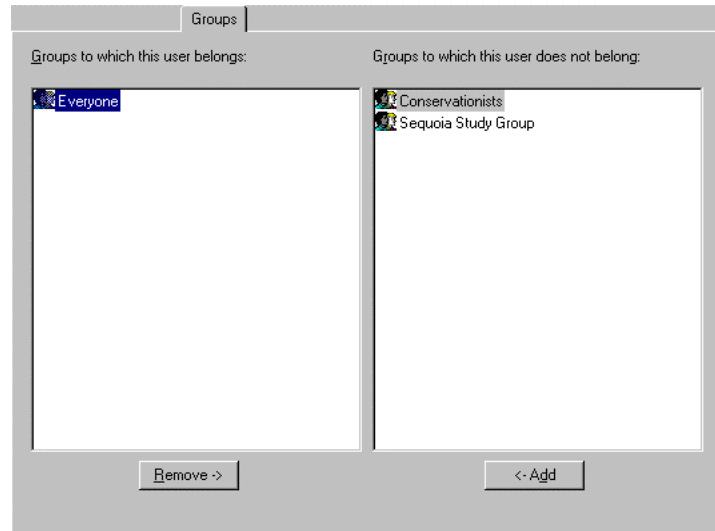
Sorting selected files according to object attributes

“Filter Selection Criteria,” Chapter 5

Groups Tab

Applicable Objects: *User Object*

Shows the groups to which the current user belongs. Use this tab to add or remove a user from a group.



Groups Tab

Groups to which this user belongs

Shows the groups to which the current user belongs. To add the user to a new group, select the group on the right side of the window and click the **Add** button; the group will move to this field. Similarly, to remove a user from a group, select the group in this field and then click the **Remove** button.

Everyone group Normally, this group will be listed in this field. When a new user is created, they are automatically added to the Everyone group. Members of this group have Modify, Delete, Create, Write, and Read permissions to the Everyone Folder. You can remove a user from this group however by selecting the Everyone group and then clicking remove.

Groups to which this user does not belong

Shows the groups to which this user does *not* belong.

For more information on...	See...
Calculating effective permissions	“Effective Permissions,” Chapter 11
Assigning members to groups	“Setting up Users,” Chapter 11

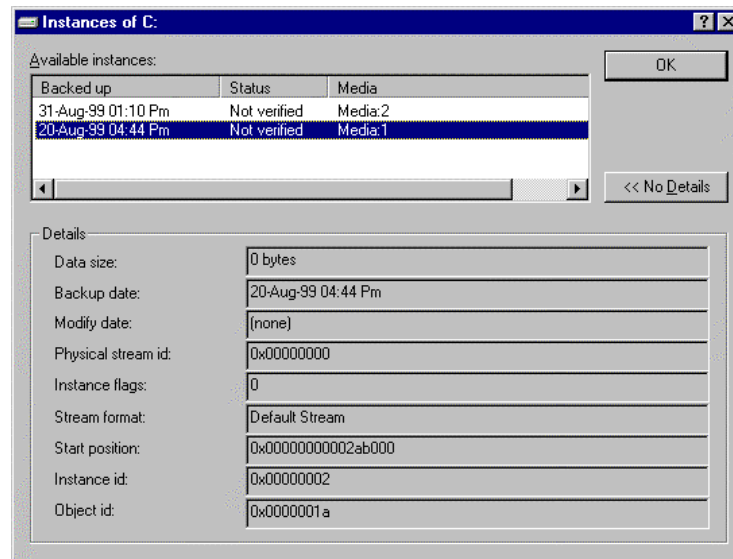
Instance Window

Opened by clicking the **Check** button on the toolbar of the **Selection** tab of either a restore or verify job. This window is used to select an instance of the object targeted on the **Selection** tab.

Each time a file is backed up an *instance* of that file is created. There may be multiple instances of files stored on different media created by different backup jobs. TapeWare keeps track of all the instances of each file in its database and the media on which each instance is stored. When media is overwritten or deleted, TapeWare deletes those instances from its database as well.

When you select a file for restoring, TapeWare initially selects the **<Latest>** instance. To select an instance other than the latest instance, use the **Instance** window.

The **<Latest>** instance is a wildcard and automatically selects the most recent instance. Which instance is selected is updated as the restore or verify job is run.



Instances Window

Available instances

Shows a list of the instances of the file and the media on which those instances are stored. The **Backed up** field shows the date and time the job was run. All files and directories backed up during a single job will be listed with the same date and time. The **Status** field shows whether or not the file was verified when

the job was run. It is either **Verify complete**, **Verify failed**, or **Not verified**. The **Media** field shows the media on which the instance is stored.

Select which instance you want to restore by highlighting it and clicking **Ok**.

Details

Shows more information about whatever instance of a file is highlighted. When you click on this button, TapeWare displays various details it uses to manage the file in its database, including the following:

Data size: Shows the size of the data fork for the selected object. For folders and directories, this figure equals 0 bytes; for files, the size of the file.

Backup date: The date and time this instance was created.

Modify date: The last time the file or directory was modified. This information is recorded from the operating system when the file is backed up.

Physical stream id: Shows internal data used by TapeWare to manage the instance.

Instance flags: Shows internal data used by TapeWare to manage the instance.

Stream format: Shows what format the instance is recorded in. The stream format is controlled by the **Advanced options...** window. You can use this information to see if you can transfer the file from one operating system to another. Stream format can only be set when creating instances, not when restoring them.

Start position: Shows internal data used by TapeWare to manage the instance.

Instance id: Shows internal data used by TapeWare to manage the instance.

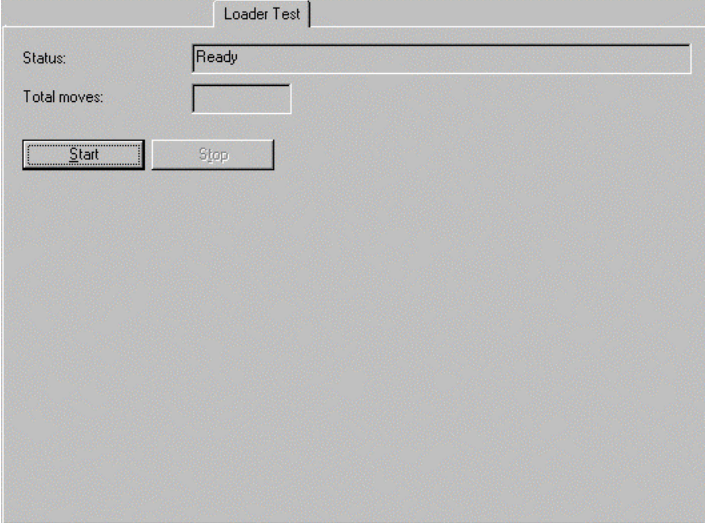
Object id: Shows internal data used by TapeWare to manage the instance.

For more information on...	See...
Selecting instances	"Selecting Instances of Files for Restore Jobs," Chapter 5
Using filters to select instances	"Backup Range," Chapter 5

Loader Test Tab

Applicable Objects: *Autoloader*

Shows the status of the current object. Performs a test of the ability of autoloader to load and unload media to the associated device. Click **Start** to begin the test.

The screenshot shows a software window titled "Loader Test". Inside the window, there is a "Status:" label followed by a text box containing the word "Ready". Below this is a "Total moves:" label followed by an empty text box. At the bottom left of the window, there are two buttons: "Start" and "Stop". The "Start" button is highlighted with a dashed border.

Loader Test Tab

Status

Shows an informational message about the current activity of the autoloader.

Total Moves

The total number of moves performed during the current test.

Logon Control Tab

Applicable Objects: *User Object*

Controls the ability of a user to log on to the current TapeWare database zone. Controls whether passwords are required, whether and when the password must be changed, whether an account has expired, and the number of connections a user can have to the network.

The screenshot shows the 'Logon Control' tab with the following settings:

- Expiration:**
 - ☐ Account is disabled
 - Date account expires: 01-Sep-99
- Grace logons:**
 - ☒ Grace logons allowed
 - Limit grace logons: 6
 - Remaining grace logons: 6
- Connection:**
 - Number of concurrent connections: 250
 - User can log on from these machines: (empty list box with Add... and Delete buttons)
- Password:**
 - ☒ Require password
 - ☒ Allow user to change password
 - Minimum password length: 5
 - ☒ Require unique passwords
 - ☒ Force periodic password changes
 - Days between forced changes: 40
 - Date password expires: 27-Sep-99

Logon Control Tab

Expiration

Controls whether the current account is expired. A user account can expire on a given date or can be disabled manually.

Date account expires: Specifies a date when the account will no longer be active. When the account expires, TapeWare disables the account and checks the **Account is Disabled** box. This user will be unable to log on until the **Account is Disabled** box is cleared.

Account is disabled: Checked when the account has expired. If manually checked, the account will be disabled.

To make a disabled account active again, clear the **Account is Disabled** check box.

Grace logons

Controls whether the current user can log on if their password has expired. If the **Force periodic password changes** box is checked, TapeWare will prompt the user to change their password as the user logs on. **Grace logons** refer to the number of times a user can log on to TapeWare with their old password when that password has expired. For example, if the **Allow Grace Logons** box is checked and the number of grace logons is set to 2, the user will be allowed to log on two times using their old password, even though that password has expired. The third attempt to log on will be denied.

Grace logons allowed: If checked, the user will be allowed to log on with an expired password.

Limit grace logons: Shows the maximum number of grace logons allowed. Can be set to desired number.

Remaining grace logons: The number of remaining grace logons. Automatically adjusted each time user logs on with expired password. Additionally, can be manually set to desired number.

Password

Controls whether a given user must have a password and the parameters of that password.

Require password: When checked, TapeWare requires the user to have a password. Note that when unchecked, if user has valid password, TapeWare will still require user's password.

Minimum password length: Determines the minimum length of the password.

Require unique passwords: If marked, TapeWare will check to see if the password is unique.

Force periodic password changes: If marked, TapeWare will force the user to change their password at the specified intervals.

Days between forced changes: Specifies the interval between forced password changes.

Date password expires: Calculated date of next password expiration. Can be manually set and adjusted.

Connection

Controls from which machines a user can log on to the current database and how many connections a user can have simultaneously.

Number of concurrent connections: Limits the number of a user's concurrent connections. Controls how many different logons a user may have simultaneously from different workstations or file servers. For example, if the **Number of concurrent connections** is set to 5, this user will be allowed to log on to TapeWare from up to five separate workstations or file servers at once.

User can log on from these machines: Controls from which machines a user can log on to the current database. The user will only be allowed to log on to TapeWare from the machines listed in this list box. To add machines, click **Add...** and select the appropriate machine from the **Browse** window. Note: if not machines are specified, the user will be able to log on from any machine.

For more information on...

Setting up new users

See...


"Setting up Users," Chapter 11

Logs Tab

Applicable Objects: *Backup Job, Restore Job, Verify Job*

Shows the available logs for the current job.

The **Log option** on the **Options** tab determines whether TapeWare keeps a log of the current job and, if so, what kind of log it keeps. After a job is run, TapeWare creates a log of the job. You can view or print this log to see which files were successfully or unsuccessfully backed up, restored, or verified.



Logs Tab

Logs for this object

Lists the available logs for this job. To view a log, click on the **View** button. TapeWare opens the log with the text editor you specified in the **Preferences...** window. You can print the log from the text editor. Before printing the log, check its length. Some logs can be quite long.

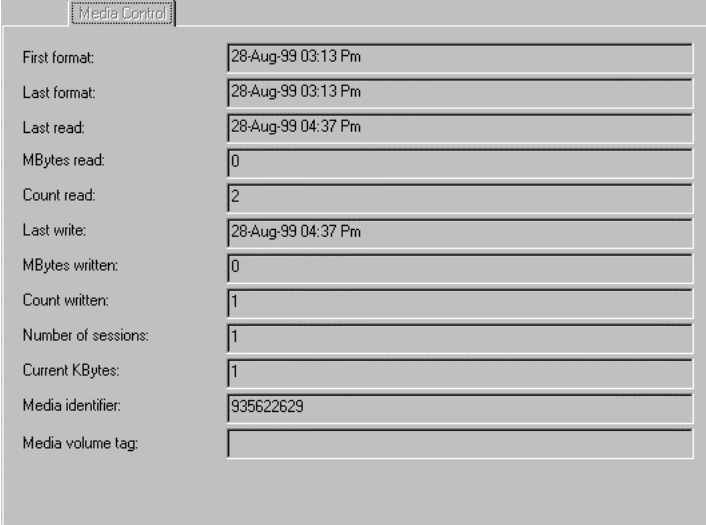
Logs can be deleted by selecting them and then clicking **Delete**. The maximum number of logs is 250. TapeWare will delete the oldest log when you reach this maximum.

For more information on...	See...
Creating logs	"Log option," Chapter 7
Specifying a text editor	"Preference Window," Chapter 12

Media Control Tab

Applicable Objects: *Media*

Shows information TapeWare stores in its database about the current valid media.



The screenshot shows a window titled "Media Control" with a list of media statistics and their values in a table-like format.

First format:	28-Aug-99 03:13 Pm
Last format:	28-Aug-99 03:13 Pm
Last read:	28-Aug-99 04:37 Pm
MBytes read:	0
Count read:	2
Last write:	28-Aug-99 04:37 Pm
MBytes written:	0
Count written:	1
Number of sessions:	1
Current KBytes:	1
Media identifier:	935622629
Media volume tag:	

Media Control Tab

First format

Shows the date and time the current media was first formatted.

Last format

Shows the date and time the current media was last formatted.

Last read

Shows the date and time the current media was last read. Matches the last time files were read for either a restore or verify job, or the verify pass of a backup job.

Mbytes read

The total number of megabytes read during either the last verify or restore job.

Count read

The total number of times this media has been read. This number is not reset when the media is overwritten.

Last write

Shows the date and time the current media was last written to. Matches the last time files were written to the media during a backup job.

Mbytes written

The total number of megabytes written during the last backup job.

Count written

The total number of times this media has been written to. This number is not reset when the media is overwritten.

Number of sessions

The number of jobs this media has been used with. This number is reset each time the media is overwritten. In general, this number is the number of backup jobs currently stored on the media, including the first overwrite job and subsequent append jobs.

Current Kbytes

The cumulative size of the data stored on the media in Kbytes. This number is reset when the tape is overwritten.

Media identifier

Unique number used internally by TapeWare to track the media in the database.

Media Volume tag

The volume tag for the physical media. In general, this number equals the bar code number on the media. Used by devices with optical readers to identify media. Only available for media used with devices that support volume tags.

For more information on...

Formatting media

See...

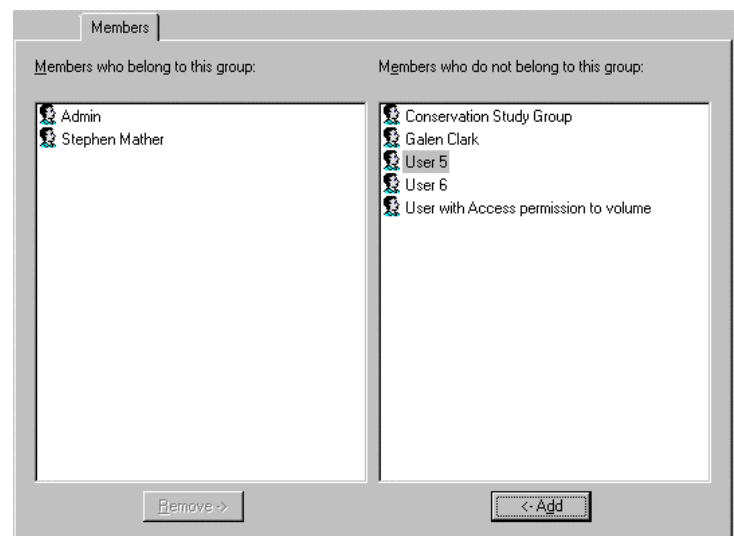
“Creating New Media,” Chapter 9

Members Tab

Applicable Objects: *Group Object*

Shows the members of a group. Adds and removes members from groups. A user's group membership is used to calculate their effective permissions.

Note that gaining permissions through group membership does not exclude gaining permissions other ways. A user has *direct permissions* to an object as a result of three situations: (1) the user is listed on the **Permissions** tab of the object (and reciprocally, the object is listed on the user's **Permissions** tab); (2) the user is equivalent to a user who has direct permissions to the object; and (3) the user is a member of a group that has direct permissions to the object. Note that these three ways of gaining permissions are not mutually exclusive: a user can have direct permissions in only one of the ways, in two of the ways, or in all three of the ways.



Members Tab

Members who belong to this group

Shows which users are a member of the current group. To add a user to the group, select the user on the right side of the window and click the **Add** button; the user will move to this field. To remove a user from a group, select the user in this field and click the **Remove** button; the user will move to the right side of the window.

Members who do not belong to this group

Shows which users are not members of the current group.

For more information on...**See...**

Adding users to groups

“Groups Tab, User Object,” Chapter 11

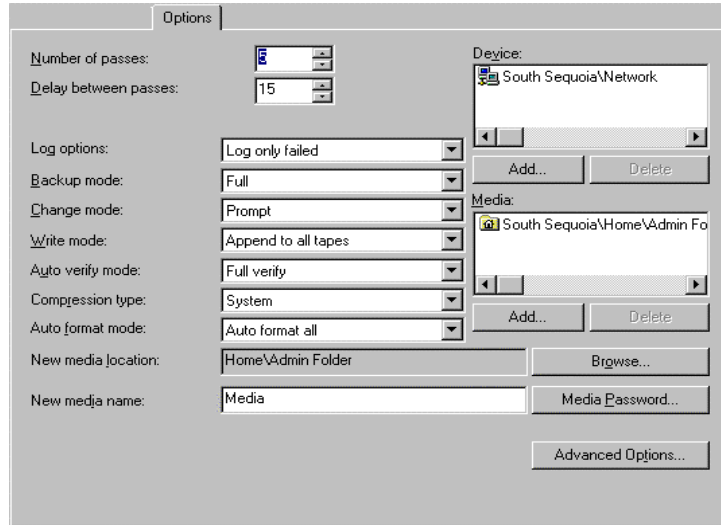
Calculating effective permissions

“Effective Permissions,” Chapter 11

Options Tab

Applicable Objects: *Backup Job, Restore Job, Verify Job*

The **Options** tab on the property sheet of a job controls various parameters important to how TapeWare runs backup, restore, and verify jobs.



Options Tab

Number of passes

Determines how many times TapeWare will attempt to access a file. The default value is **5**. Applies to backup, restore, and verify jobs.

If TapeWare needs to use a file and that file is already open, that is, it is currently being used by another user, TapeWare will attempt to wait until that user is through using the file. Each time TapeWare attempts to open a file is called a **pass**.

For backup jobs, when TapeWare is unable to back up a file on the first pass, it attempts to back up that file on subsequent passes. On the last pass, TapeWare opens the file in *shared* mode.

Delay between passes

Determines how many seconds TapeWare waits before attempting the next pass. If your historical usage suggests that many files are being opened on the last pass, consider increasing this parameter to a higher number.

Log options

Determines whether TapeWare keeps a log of the current job and, if so, what kind of log it keeps. The **Log option** is either **None**, **Log only failed**, **Log only completed**, or **Log all**.

After a job is run, you can view or print the log to see which files were successfully or unsuccessfully backed up, restored, or verified. The default value is Log only failed, which writes in the log any files which were not successfully backed up, restored, or verified.

None Instructs TapeWare to not keep a log of the backup job as it runs.

Log only failed Instructs TapeWare to log the name of any file selected but, for some reason, not successfully backed up, restored, or verified.

Log only completed Instructs TapeWare to log the name of any file selected for and successfully backed up, restored, or verified.

Log all Instructs TapeWare to log the name of every file selected and whether or not that file was successfully backed up, restored, or verified.

Backup mode

Applicable to backup jobs only. Determines whether all files or only changed files are backed up. For scheduled automatic rotation jobs, TapeWare uses the backup mode Type parameter from the **Schedule** tab; for unscheduled or manual jobs, TapeWare uses the parameters specified in this list box.

The **Backup mode** is either **Full**, **Incremental**, **Differential**, or **Snapshot**.

Full This parameter instructs TapeWare to back up all selected files. TapeWare does not check to see if the file's archive bit has been set. After backing up each file, the archive bit is unchecked.

Differential This parameter instructs TapeWare to back up all selected files that have changed since the last *full* backup. TapeWare selects only files with their archive bit checked. After restoring them, the archive bit is not changed (it is left checked). These files will be selected for the next backup job, whether full, differential, or incremental.

Incremental This parameter instructs TapeWare to back up all selected files that have changed since the *last* backup. TapeWare selects only files with their archive bit checked. After restoring them, the archive bit is changed to unchecked. These files will not be selected for the next differential or incremental job unless they change.

Snapshot This parameter instructs TapeWare to back up all selected files. TapeWare does not check to see if the file's archive bit has been set. After backing up each file, the archive bit is *left unchanged*.

Change mode

Determines what action TapeWare will perform when it fails to find the media it was expecting to use for a job. When TapeWare runs a job, if the job uses specific media, TapeWare scans the network for devices with that media. If it does not find the media it expects, its response is determined by the **Change mode** parameter.

Skip device This parameter instructs TapeWare to skip the designated device and look for other devices on the network with the proper media. TapeWare continues to scan for the correct media until it is found. If all of the available devices contain the incorrect media, the job will be terminated.

Force to append This parameter instructs TapeWare to append data to whatever media it finds in the designated backup device. If it cannot find the correct media, TapeWare appends data to whatever media is available. This option will insure that the job runs, if the media contains enough room to complete the job.

Prompt This parameter instructs TapeWare to continue to scan for the expected media and to send an alert warning that the proper media has not been found. This option will not allow a job to run with any other media except with the expected media. Additionally, this option will not search for another device that might contain the proper media.

Write mode

Determines whether the old data on the media is *overwritten* with new data, or whether the new data is *appended* to the end of the old data. When media is overwritten, all of the data previously stored on it is lost. Appending data will preserve the old data.

Applies only to backup jobs. For scheduled automatic rotation jobs, TapeWare defaults to **Overwrite all** mode; for unscheduled and manual jobs, TapeWare uses the parameters specified in this list box.

The **Write mode** is either **Append to all**, **Append to first overwrite others**, or **Overwrite all**.

Append to all Instructs TapeWare to append all data to the end of the media. No data is overwritten. Select this parameter for permanent storage.

Append to first, overwrite others Instructs TapeWare to append data to the end of the first media, but to overwrite all media that follows. For example, TapeWare

will not overwrite the first tape inserted, but will overwrite the second, third, and later tapes. This parameter is useful if you have a set of media with old data you no longer need. By selecting this option, TapeWare preserves your most recent data on the first media, but overwrites older, unneeded media.

Overwrite all Instructs TapeWare to overwrite all media. All data on media that is overwritten is lost. Use this option for tapes that are going to be recycled.

Auto verify mode

Verifies if a file has been correctly written to tape. Compares file with original on volume. Applies to both backup jobs and verify jobs.

The **Auto verify mode** is either **Full verify**, **No verify**, or **Quick verify**.

Full verify Instructs TapeWare to compare every selected file stored on media with the original file from the workstation or file server. Checks to see if file can be read and then checks to see if instance matches original. This default value is strongly recommended.

Quick verify Instructs TapeWare to read every selected file on the tape and to verify that the instance is in readable condition. It does not check to see that file instance stored on the media matches the original file, only that the data stored on the tape (incorrect or not) can be read. While selecting this option can save time, it is nonetheless not recommended.

No verify Instructs TapeWare to skip the verification step. It is not recommended.

Software Compression

The Software Compression mode controls how TapeWare compresses or maintains the compression of files and directories.

None This parameter instructs TapeWare to write all data to the tape in a decompressed format. If the file is stored on disk in a compressed format, the file will be decompressed before writing. This option is useful if the device supports hardware data compression and the files are to be restored to a different operating system.

Standard This parameter instructs TapeWare to write all data to tape in the TapeWare compression format. If the file is stored on disk in a compressed format, the file will be decompressed before being re-compressed by TapeWare. This options is useful if the tape device does not support hardware data compression and the files are to be restored to a different operating system.

System This parameter instructs TapeWare to write all data to tape in the same mode it is stored on disk. If the file is stored on disk in compressed format, TapeWare will write the data in the hosts compressed format. If the file is no

compressed on disk, TapeWare will store the file on tape in a non-compressed format. This option is useful if the hardware does support data compression and the files are to be restored to the same operating system.

Both This parameter instructs TapeWare to write all data stored on disk in compressed format, but those files that are not compressed on disk will be stored in the TapeWare compression format. This option is useful if the hardware does not support data compression and the files are to be restored to the same operating system.

Auto format mode

Determines whether or not TapeWare will format media automatically. Applies only to backup jobs.

Before data can be written to media, the media must be formatted. When media is formatted, any data on it is lost. Tapes and other media are formatted when TapeWare does not recognize the media, that is, when it has no information in its database about that particular media. This will occur when the tape is blank, it has been erased, it is first used, or it has been deleted from the database. If you want to use media but do not wish to format it, you must import the media.

The **Auto format mode** is either **No auto format**, **Auto format all**, or **Auto format blank**.

No auto format When selected, if TapeWare encounters media that needs to be formatted (either blank or unrecognized media), it sends an alert to the alert window. While waiting for a user reply, TapeWare scans the network for devices with the media it was expecting. When this parameter is selected, TapeWare will wait until the user replies to the alert before formatting unrecognized media.

Auto format all Instructs TapeWare to automatically format all of the media inserted into the tape drive which require formatting. With this parameter selected, TapeWare will automatically format all new (or blank) tapes and all unrecognized media.

Auto format blank: Instructs TapeWare to automatically format all new or blank media. However, if TapeWare encounters unrecognized media, it sends an alert to the alert window and then scans the network for the media it was expecting. This parameter can help prevent data from being accidentally destroyed by formatting, while not needlessly querying the user before formatting a blank tape.

New media location

Specifies the folder in which TapeWare will store any new tapes created while the job is run. Applies only to backup jobs.

By default, TapeWare stores media in the current User/Group folder. Select the folder in which you wish to store any new media or tapes by clicking on the **Browse...** button and then selecting the folder from the **Browse** window. If you do not already have folders set up in which to store the media, use the **Media** tab to first create additional folders.

When TapeWare runs a scheduled automatic rotation job, it automatically creates folders. The folders are organized by the name of the job and the various rotation sets in that job. There is no reason to create these folders manually. TapeWare will automatically create these folders for you.

New media name

Specifies the name TapeWare gives to any new media it creates while running the job. Applies to backup jobs only.

For scheduled automatic rotation jobs, TapeWare automatically updates this parameter to match the media's place in the rotation schedule. For example, if the media is the first media in the yearly rotation set, TapeWare names it "Year Set 1:1".

For manual rotation and unscheduled jobs, TapeWare assigns any new media it creates the name listed in this field. This is also true for automatic rotation jobs that are "forced" to run.

Device

Specifies which tape drive or other removable media device TapeWare will use to run the backup job, restore job, or verify job.

By default, TapeWare sets this parameter to the network container. When running the job, TapeWare will use whatever device it finds on the network. If there is only one device in your database zone or if you only have permissions to one device, there is no reason to change this parameter. If there is more than one device in the current database zone, TapeWare will attempt to use all available devices while running a job.

If there are several devices on your network and you need to select a specific device to be used, specify which device the job should use by selecting it from the **Device** list. (If a machine has only one device, you need not select the device, only the machine.)

If you wish to use a device that is not shown in the **Device** list, click the **Add...** button and select the new device from the **Browse** window.

Media

Specifies the database folder in which the tape or other removable media are stored in the database. TapeWare will look here for media to use with this job. Applies only to backup jobs.

The default folder is the current User/Group folder. If you wish to use media from another folder, specify which folder by selecting it from the **Media** list.

If you wish to use a folder that is not shown in the **Media** list, click the **Add...** button and select the new folder from the **Browse** window.

Media Password...

Opens window to assign media a password. Applies only to backup jobs that format media.

When a job creates new media, you can assign that media a password. A password prevents the media from being imported into another TapeWare database and can be an important part of your overall security plan.

To assign a password to any new media your job creates, click the **Media Password...** button and type and confirm your password.

Note that passwords can only be assigned when media is formatted. Additionally, media passwords are only required when importing media.

For more information on...	See...
Setting job options	Chapter 7, “Job Options”
Backup mode	“Backup Mode,” Chapter 6
Formatting media	“Creating New Media,” Chapter 9

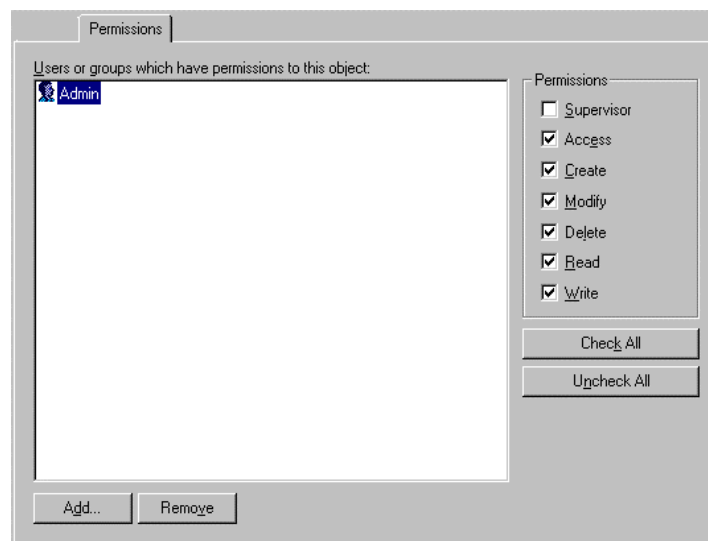
Permissions Tab

Applicable Objects: *All database objects*

For user objects and group objects, lists the objects to which the current user or group has permissions. For all other objects, shows the users or groups which have permissions to the current object.

Use this tab to grant users or groups permissions to objects. Note that permissions can be granted from either the property sheet of the database object or the property sheet of the user or group. Either way, the permissions appear on the appropriate corresponding object's **Permissions** tab. For example, if Galen Clark is granted permissions to the **C:** volume from the **Permissions** tab on his property sheet, the **Permissions** tab on the property sheet of the **C:** volume will list Clark as a user who has permissions. Alternatively, if Clark is granted permissions from the property sheet of the **C:** volume, the appropriate permissions will appear on Clark's **Permissions** tab.

Note additionally that a user has direct permissions only to those objects listed on that user's **Permissions** tab. Any and all other effective permissions to other objects are calculated through inherited permissions, through equivalencies, or through groups.



Permissions Tab

Users or groups which have permissions to this object

Lists the users or groups which have permissions to the current object. (Not applicable to either group objects or user objects.)

To see what permissions each user or group has, highlight the user or group; that group’s or user’s permissions will be displayed in the **Permissions** field.

To add a user or group, click the **Add...** button. To remove a user or group, highlight the group or user and then click the **Remove** button.

Objects to which this user or group has permissions

Lists the objects to which the current user or group has permissions. (Applicable to only groups objects and user objects.)

To see what permissions each user or group has to a particular object, highlight the object; that group’s or user’s permissions to that object will be displayed in the **Permissions** field.

To grant the user or group permissions to a new object, click the **Add...** button and select the appropriate permissions. To end permissions to an object, highlight the object and then click the **Remove** button.

Permissions

Shows the permissions granted to the currently highlighted object, user, or group. Check or uncheck the appropriate boxes to grant or restrict permissions.

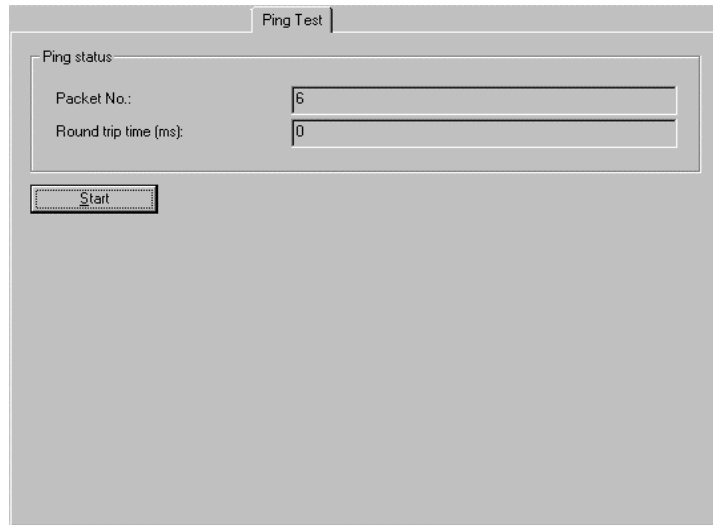
For more information on...	See...
Calculating effective permissions	“Effective Permissions,” Chapter 11
Assigning users or groups permissions	“Permission’s Tab...,” Chapter 11
Differing types of permissions	“Permissions Reference,” Chapter 11

Ping Test Tab

Applicable Objects: *Machine*

Used to measure how long it takes an echo packet to travel back and forth between two machines on the TapeWare network.

Sent between two machines running TapeWare. Sent from machine TapeWare is currently running on and machine whose property sheet is open.

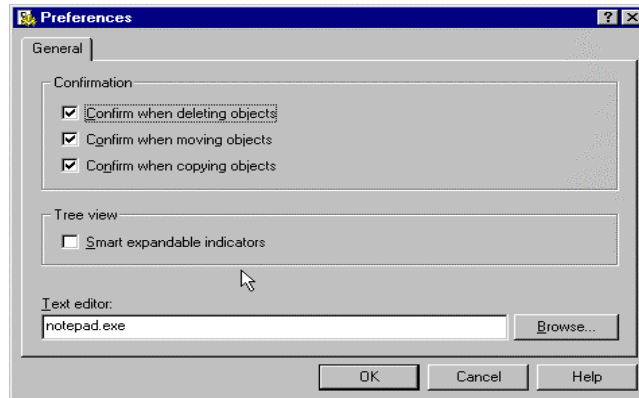


The screenshot shows a window titled "Ping Test" with a tab labeled "Ping Test". Inside the window, there is a section titled "Ping status" containing two input fields: "Packet No." with the value "6" and "Round trip time (ms)" with the value "0". Below these fields is a "Start" button.

Ping Test Tab

Preferences Window

Sets options controlling user interface.



Preferences Window

Confirmation

Allows user to determine if TapeWare will prompt before completing certain commands.

Confirm when deleting objects: When checked, TapeWare will open **Confirm Delete** window before deleting objects. Requires additional user input before deleting database objects.

Confirm when moving objects: When checked, TapeWare will open **Confirm Move** window before moving objects. Requires additional user input before moving database objects.

Confirm when copying objects: When checked, TapeWare will open **Confirm Copy** window before pasting copied objects. Requires additional user input before duplicating database objects.

Tree View

Determines whether TapeWare calculates if containers in hierarchical, tree views have containers within them. Determines whether or not TapeWare displays expand tree icons next to containers that cannot be expanded (i.e., they do not contain additional containers within them.)

Smart expandable indicators: When checked, TapeWare looks in every container in a hierarchical, tree view to see if there are containers within it. If not, the expand tree icon is not displayed.



The Expand
Tree icon

When unchecked, TapeWare displays the expand tree icon for every container. In this case, TapeWare only checks to see if the container holds other containers when the expand tree icon is clicked. Clearing this option generally results in faster display of the hierarchical tree, although with erroneous expand tree icons.

Text editor

Determines what text editor (word processor) TapeWare will use when displaying logs. When you select a job log on the **Logs** tab and click **View**, TapeWare opens the log using the text editor specified in this field. You can change the text editor by clicking the **Browse** button and selecting another text editor. Note that the default Windows text editor is Notepad.

Schedule Tab

Applicable Objects: *Backup Job, Verify Job, Restore Job*

Controls when and how often a job is run.

Schedule Tab

Type

Sets the type of schedule.

For restore and verify jobs, the type is either **Not Scheduled** or **Manual**. In addition, for backup jobs, there are eight automatic rotation schedules plus the custom schedule.

Not Scheduled: The job will run only when instructed to do so. Uses parameters set on job's **Options** tab.

Manual: Turns on scheduling calendar, allowing job to be scheduled to run repeatedly. Job will run every day indicated on calendar. Uses parameters set on job's **Options** tab. User manually controls set count, media rotation, media name, and backup mode.

Custom: Turns on automatic rotation schedule. Applies to backup jobs only. Job will run every day as scheduled in the calendar. TapeWare will automatically update the **Backup mode**, **Write mode**, and **New media name** parameters on the **Options** tab of the job when it runs the job as scheduled. (These parameters are not updated if the job is manually "forced" to run by a user.) Allows the user to determine the set count for each set type; however, TapeWare will

automatically control the implementation of these features. When first selected, initially defaults to GFS-25 schedule.

Automatic Rotation Schedules: There are eight automatic rotation schedules: GFS 30-tape, GFS 25-tape, GFS 20-tape, Simple 12-tape, Simple 11-tape, Simple 10-tape, Simple 6-tape, Simple 4-tape. Applies only to backup jobs. Job will run every day scheduled in calendar. TapeWare will automatically update the **Backup mode**, **Write mode**, and **New media name** parameters on the **Options** tab of the job when it runs the job as scheduled. (These parameters are not updated if the job is manually “forced” to run by a user.) Set count for each set type is predetermined; TapeWare automatically controls media rotation.

Start time

Specifies the time of day the job will run. Note that jobs can run concurrently.

Daily, Weekly, Monthly, Yearly

Indicates the backup mode and set count for the **Daily**, **Weekly**, **Monthly**, and **Yearly** media sets respectively. Backup mode (**Type**) and set count (**Count**) can be set by the user when the **Custom** schedule is selected.

End of Week

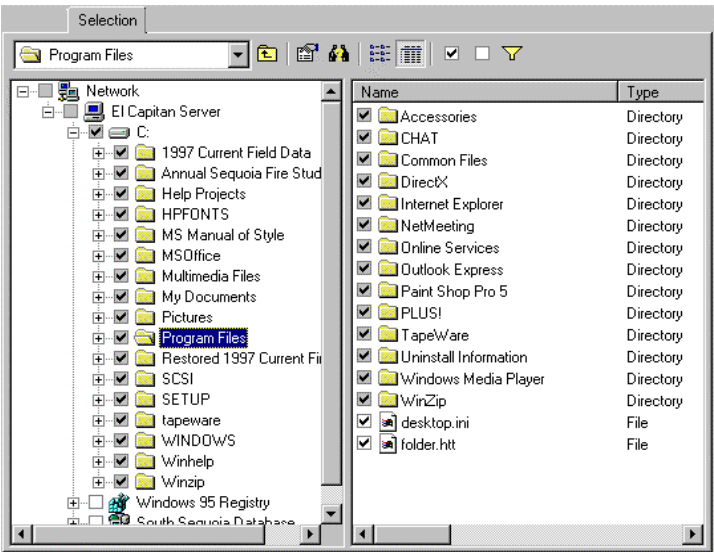
Indicates the day of the week that TapeWare will use to schedule **Weekly** backup jobs. Change the list box to match whatever day of the week TapeWare should run weekly jobs.

For more information on...	See...
Differences between schedule types	“Which Built-in Schedule to Select,” Chapter 6
Manually forcing scheduled jobs to run	“Forcing Scheduled Jobs to Run,” Chapter 8
Scheduling jobs to run once only	“Scheduling a Job to Run Once Only,” Chapter 10
Working with the Calendar	“Modifying Built-in Schedules,” Chapter 6

Selection Tab

Applicable Objects: *Backup Job, Restore Job, and Verify Job*

Specifies which files or instances are selected for use by the current job.



Selection Tab

For all job types, selected files are indicated with a check mark.

Container objects (folders, volumes, machines, and networks) are grayed if they potentially contain objects within them that are selected. A container may be grayed even when no objects within it are selected. The gray indicates that if new objects are created in that container which match the filter selection criteria, then they will be selected. A container that is unmarked and not grayed is not selected and has no objects within it that are selected. A grayed check mark indicates that the container itself is selected; it may or may not contain files within it that are selected.

Files are selected for backup in three steps. In the first step, the appropriate files are selected by marking them with a check. In the second step, these marked files are sorted through using filer selection criteria. This step is optional. In the last step, as the job is run, TapeWare checks to see if it will backup all the files or only those files that have changed since the last backup job.

For more information on...	See...
Filtering and Selecting files for jobs	Chapter 5, "Selecting Files and Instances"

Selection Filters Window

Applicable Objects: *Backup Job, Restore Job, Verify Job*

Used to sort files for jobs. Excludes or “filters out” files that do not meet the specified selection criteria.

The selection filters are applied to all of the volumes, folders, and files that have been marked for backup. *You cannot apply different filters to different folders or volumes.* If a folder or other container has been marked for backup, TapeWare uses the selection filters to sort through the files and unmark any files that do not meet the selection criteria. *TapeWare does not use the selection filters to add files to the backup set.*

The filter criteria are applied to all marked files, regardless if they were marked before or after the filter criteria were specified. After specifying the selection filter criteria, you can then mark or unmark files, folders, and volumes for backup. Additionally, you can change the filter criteria at any time; TapeWare will automatically reapply the new selection filter criteria to the marked folders and files.

Selection Filters

Backup range: <Any Date> >>

Modify range: <Any Date> >>

Create range: <Any Date> >>

Access range: <Any Date> >>

Size range: <Any Size> >>

Instance range: <Any number of instances> >>

Wildcard type: *.*

Must match:

Cannot match:

Media:

Required attributes:

- ☐ Read Only
- ☐ Hidden
- ☐ System
- ☐ Execute Only

Exclude attributes:

- ☐ Read Only
- ☐ Hidden
- ☐ System
- ☐ Execute Only

☒ Parents

☒ Children

Reset All Add... Delete OK Cancel

Selection Filters

Backup range

Selects files according to their backup date. The backup date is assigned to a file each time it is backed up. The backup date for a file is the same as the *last* time it was backed up.

Modify range

Selects files according to their modify date. Each time a file is modified, its modified date is updated. You can use this filter to select files whose modified date matches your criteria. TapeWare checks the directory information on the volume to see if the file should be included in the job.

Create range

Selects files according to their create date. When a file is first created, it is assigned a create date. You can use this filter to select only those files which match your criteria. TapeWare checks the created date for each file stored in the directory of the volume and uses this to sort files for the job.

Delete range

Applies only to restore jobs. Selects files according to their delete date.

When files have been backed up and are later deleted, TapeWare marks the file as having been deleted and assigns it a delete date. This filter instructs TapeWare to only restore files which have a delete date that matches the selection criteria. If a file has not been deleted, it will not have a delete date and will not be selected.

Access range

Selects files according to their access date. Each time a file is read, whether or not it is modified, its access date is updated by the operating system. You can use this information to select and filter files.

Size range

Selects files according to their size.

Instance range

Selects files according to their instance date. Each time TapeWare backs up a file, it creates a new instance of that file and assigns it an instance date. For restore and verify jobs, this is the only filter that selects *instances* of files, rather than *files* themselves.

Wildcard type

Specifies the wildcard format used by the **Must match** and **Cannot match** filters. One of three types of wildcard formats are available: DOS, Long, or UNIX.

Must match

Selects files that match specified wildcards. Only files that match the wildcard indicated in this field are selected.

Specify multiple wildcards by separating each with a semicolon, “;”. TapeWare selects any file that matches any one of the wildcards. For example, if you enter “*.exe;*.doc” in the **Must match** field, TapeWare selects all files that *either* have the .exe extension *or* the .doc extension.

Cannot match

Unselects files that match the specified wildcard. Excludes any files that match the wildcards. You can specify multiple wildcards by separating them with a semicolon; if you specify multiple wildcards, TapeWare excludes any file that matches any one of the wildcards you specify.

Required attributes

Selects files according to attributes controlled by the operation system.

Operating systems track certain features of files called attributes that they use to manage these files. In this field, if an attribute is checked, TapeWare only selects those files which have these attributes.

You can select multiple attributes. In this case, TapeWare only selects those files that meet *all* of the required attributes.

Note that some of these attributes are only supported by certain operating systems. If you specify an attribute that is specific to a particular operating system, then only files created under that operating system will be selected for backup.

Exclude attributes

Unselects files according to attributes controlled by the operating system.

This field works like the **Required attributes** field except that TapeWare excludes files that match these attributes.

You can select multiple attributes. TapeWare excludes any file that has *any one* of the attributes. For example, if you select the **Hidden** and **System** attributes, a file will be excluded if it has *either* the **Hidden** attribute *or* the **System** attribute.

Parents

Determines whether or not the directory information is selected.

When this option is checked, when TapeWare backs up or restores a file, it also backs up or restores the directory information relevant to that file. This option must be checked in order for folders and other directory data to be backed up or restored. When this option is not checked, TapeWare will not back up or restore any parent information for any file backed up. If unchecked, directory information about folders and volumes is not backed up.

Children

Determines whether or not the files are selected .

When this option is checked, TapeWare backs up and restores the selected files. If you want only to back up or restored the marked *directories*, uncheck this option. When the **Children** box is unchecked and the **Parents** box is checked, TapeWare backs up the directory structure, but not the files stored in the directories (that is, in the folders).

Media

Selects files which have a valid instance on the media listed in this field.

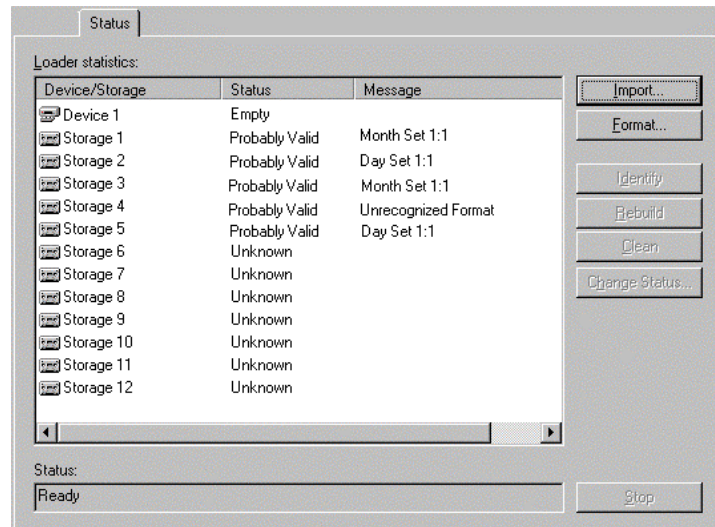
TapeWare tracks instances of files and the media on which those instances are stored. You can use this information to sort files according to the media on which they appear. Only files with instances on the media in the **Media** field will be selected for a job. If there are multiple media shown in the **Media** filter field, only files which have a valid instance on *all* the media listed will be selected.

For more information on...	See...
Using file selection filters	"File Selection Criteria," Chapter 5
Selecting deleted files only	"Selecting Deleted Files for Restoring," Chapter 10
Using filters to select instances	"Backup Range," Chapter 5
Using filters to copy directories	"Copying a Directory Structure," Chapter 10

Status Tab

Applicable Objects: *Autoloader*

Shows the status of the Device and Media associated with an autoloader. Allows user to select storage slots and media and perform various functions on the media in the autoloader.



Status Tab

Loader statistics

Shows information about the current magazine in the autoloader.

Device/Storage Specifies the name of either the device or the storage slot.

Status Shows the current or likely status of the device or storage slot. If **Valid**, the slot is known to hold media that is in the current database. If **Invalid**, the slot holds media that is definitely known not to be in the current database. If **Probably Valid**, that slot held valid media previously, but TapeWare will check to see that the media is valid before using it. When you exit and restart TapeWare, media marked **Valid** is reset to **Probably Valid**. If **Unknown**, the status of the slot is not known, normally because this slot has not yet been used. If **Clean**, TapeWare assumes the storage slot holds a cleaning cartridge.

Message Displays the name of **Valid** or **Probably Valid** media.

Import...

Opens the **Import Media** window for the selected storage slots.

Format...

Opens the **Format Media** window for the selected storage slots.

Identify

Identifies the selected media. Reads the media header information and checks to see if media is in current database.

Rebuild

Identifies the selected media. If autoloader supports optical scanning of media bar codes, uses bar codes to identify the media, comparing it to the current database. If autoloader does not support optical scanning of media, reads the media header information and checks to see if media is in current database.

Clean

Performs a cleaning cycle on the selected device. Requires that a cleaning cartridge be inserted into storage slot and a slot be marked **Clean**.

Change Status...

Opens the **Change Status** window. Lets user assign media in selected storage slot either **Unknown**, **Empty**, or **Clean** status. If **Clean**, TapeWare will use the media in this slot when performing a cleaning cycle. If **Unknown**, TapeWare will identify the media in the slot before using it.

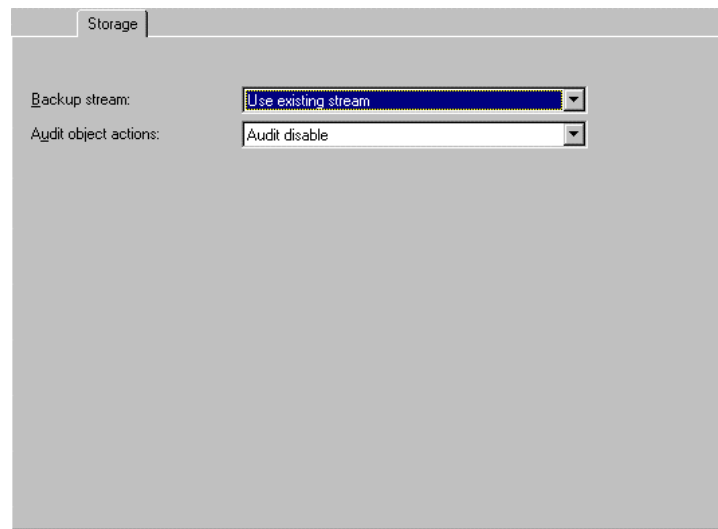
Storage Tab

Applicable Objects: *File, Directory, Volume*

Determines whether or not a unique data stream is created for the current object when a backup job is run.

TapeWare is capable of controlling up to 16 data streams simultaneously. Data streams are automatically created for each machine object; this parameter cannot be changed.

By default, new data streams are created for each volume, while files and directories by default use the data stream of their parent volume. For files, directories, and volumes, these parameter can be modified.



Storage Tab

Backup stream

Modifies whether or not a new data stream is created for the current object.

Use existing stream TapeWare does not create a new stream for this object. This is the default value for directories and files.

Select this parameter for a volume when you do not want to create a new stream for that volume. For example, you may wish to not create an additional stream when the volume is only a logical partition—not a physically separate device.

Create new stream TapeWare will create a new stream for this object when running a backup job. This is the default value for volumes.

Select this parameter for directories and files when you want to create a new stream for these objects. For example, to speed up a backup job, you may wish to create an additional stream for a very large file or for a RAID device.

Audit object actions

Enables or disables the audit log for this object.

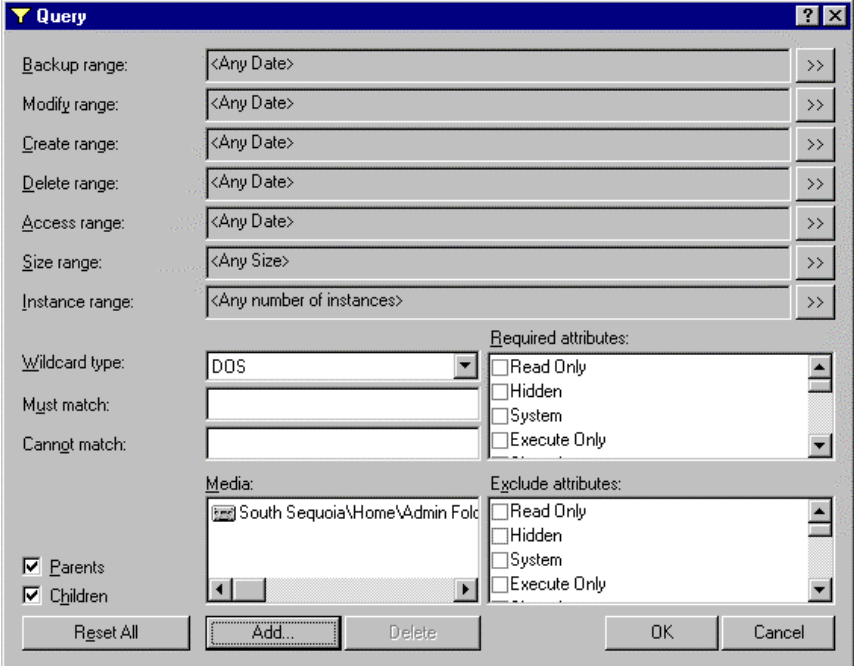
For more information on...	See...
Optimizing backup streams	“Strategies for Faster Jobs,” Chapter 10
Audit logs	“Viewing and Printing Audit Logs,” Chapter 8

Query Window

Applicable Objects: *All storage management database objects.*

Used to sort files for displaying on **Database** tab. Excludes or “filters out” files that do not meet the specified selection criteria. Displays only those objects which meet filter criteria.

The selection filters are applied to all of the volumes, folders, and files ordinarily displayed on the **Database** tab. *You cannot apply different filters to different machines or volumes.* TapeWare uses the selection filters to sort through the files and displays only those files that meet the selection criteria.

The image shows a 'Query' dialog box with a blue title bar. It contains several filter categories: 'Backup range:', 'Modify range:', 'Create range:', 'Delete range:', 'Access range:', 'Size range:', and 'Instance range:', each with a text field and a '>>' button. Below these are 'Wildcard type:' (a dropdown menu set to 'DOS'), 'Must match:', and 'Cannot match:' (text fields). To the right are 'Required attributes:' and 'Exclude attributes:' sections, each with a list of checkboxes: 'Read Only', 'Hidden', 'System', and 'Execute Only'. At the bottom left are checkboxes for 'Parents' and 'Children', both checked. A 'Media:' section shows a list box with 'South Sequoia\Home\Admin Folders'. At the bottom are buttons for 'Reset All', 'Add...', 'Delete', 'OK', and 'Cancel'.

Query Filters

Backup range

Displays files according to their backup date. The backup date is assigned to a file each time it is backed up. The backup date for a file is the same as the *last* time it was backed up.

Modify range

Displays files according to their modify date. Each time a file is modified, its modified date is updated. You can use this filter to display files whose modified date matches your criteria. TapeWare checks the directory information on the volume to see if the file should be included in the job.

Create range

Displays files according to their create date. When a file is first created, it is assigned a create date. You can use this filter to displays only those files which match your criteria. TapeWare checks the created date for each file stored in the directory of the volume and uses this to sort files.

Delete range

Displays files according to their delete date.

When files have been backed up and are later deleted, TapeWare marks the file as having been deleted and assigns it a delete date. This filter instructs TapeWare to display only files which have a delete date that matches the selection criteria. If a file has not been deleted, it will not have a delete date and will not be displayed.

Access range

Displays files according to their access date. Each time a file is read, whether or not it is modified, its access date is updated by the operating system. You can use this information to select and filter files.

Size range

Displays files according to their size.

Instance range

Displays files according to their instance date. Each time TapeWare backs up a file, it creates a new instance of that file and assigns it an instance date.

Wildcard type

Displays the wildcard format used by the **Must match** and **Cannot match** filters. One of three types of wildcard formats are available: DOS, Long, or UNIX.

Must match

Displays files that match specified wildcards. Only files that match the wildcard indicated in this field are selected.

Specify multiple wildcards by separating each with a semicolon, “;”. TapeWare displays any file that matches any one of the wildcards. For example, if you enter “*.exe;*.doc” in the **Must match** field, TapeWare displays all files that *either* have the .exe extension *or* the .doc extension.

Cannot match

Files that match the specified wildcard are not displayed. Excludes any files that match the wildcards. You can specify multiple wildcards by separating them with a semicolon; if you specify multiple wildcards, TapeWare excludes any file that matches any one of the wildcards you specify.

Required attributes

Displays files according to attributes controlled by the operation system.

Operating systems track certain features of files called attributes that they use to manage these files. In this field, if an attribute is checked, TapeWare only displays those files which have these attributes.

You can select multiple attributes. In this case, TapeWare only displays those files that meet *all* of the required attributes.

Note that some of these attributes are only supported by certain operating systems. If you specify an attribute that is specific to a particular operating system, then only files created under that operating system will be displayed.

Exclude attributes

Files with specified operating system attributes are not displayed.

This field works like the **Required attributes** field except that TapeWare excludes files that match these attributes.

You can select multiple attributes. TapeWare excludes any file that has *any one* of the attributes. For example, if you the **Hidden** and **System** attributes, a file will be excluded if it has *either* the **Hidden** attribute *or* the **System** attribute.

Parents

Determines whether or not the directories are displayed.

When this option is checked, TapeWare displays the directories for any object that meets the other display criteria.

Children

When this option is checked, TapeWare backs up and restores the selected files. If you want only to back up or restored the marked *directories*, uncheck this option. When the **Children** box is unchecked and the **Parents** box is checked, TapeWare backs up the directory structure, but not the files stored in the directories (that is, in the folders).

Media

Displays files which have a valid instance on the media listed in this field.

TapeWare tracks instances of files and the media on which those instances are stored. You can use this information to sort files according to the media on which they appear. Only files with instances on the media in the **Media** field will be displayed. If there are multiple media shown in the **Media** filter field, only files which have a valid instance on *all* the media listed will be displayed.

Configuring NetWare BTrieve Database Sets

NetWare BTrieve is a network database record manager supplied by Novell. If you are using BTrieve, you must configure a special control file to backup and restore your database.

The use of a control file is necessary because related databases (record sets that work with each other) need to be backed up as a single operation. This insures database integrity by not recording changes made to one database that are not recorded in a corresponding, related database.

You must manually specify which files are related, that is, which files need to be backed up and restored as a single operation. To do so, you create a *set*, that is, a group of database files that are related. TapeWare will open and close these files at the same time, insuring that none of the files are modified while one of the files is being backed up or restored.

When properly configured, the **Selection** tab of a job's property sheet and the **Database** tab will display a special folder named **BTrieve Database Sets**. When you expand this folder, you will see the sets you defined displayed. (The sets appear as a file with the name you gave to the set). You can select these sets just like any other file or folder for backup, restore, or verifying. Remember, however, that when you select a set, you are actually selecting the group of files defined by that set.

Defining BTrieve Sets

To define a BTrieve database set and specify which record files belong to that set, you must create and edit a special control file named TW\$BTRV.DAT. To do so, follow these steps.

First, create or edit the SYS:\SYSTEM\TW\$BTRV.DAT file. Follow this example:

```
:LOAD EDIT SYS:\SYSTEM\TW$BTRV.DAT
```

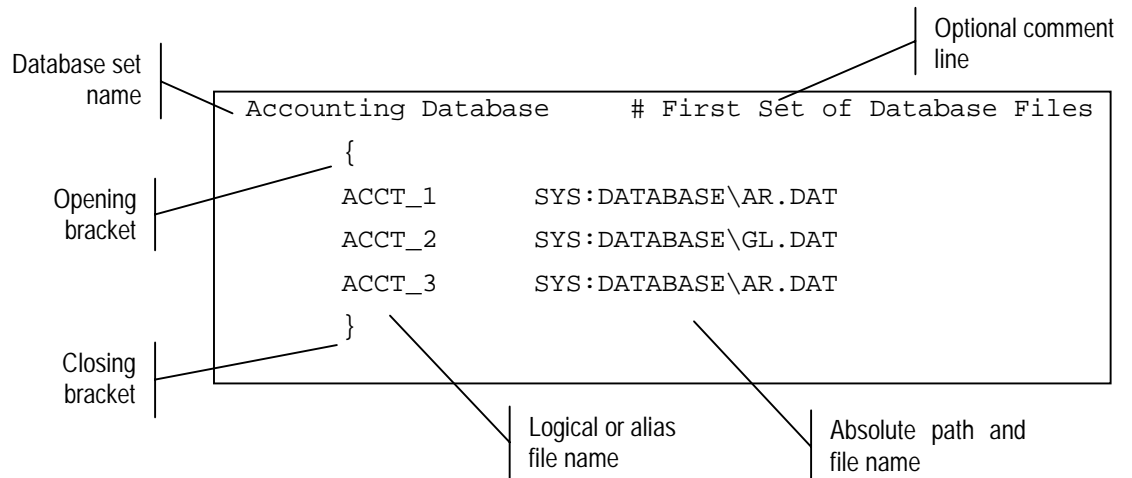
Then specify the contents of the TW\$BTRV.DAT file. Follow this format carefully.

```
Accounting Database           # First Set of Database Files
{
    ACCT_1      SYS:DATABASE\AR.DAT
    ACCT_2      SYS:DATABASE\GL.DAT
    ACCT_3      SYS:DATABASE\AR.DAT
}

Accounting Database Temp      # Redirected accounting files
{
    ACCT_1      SYS:DATABASE\OLD\AR.DAT
    ACCT_2      SYS:DATABASE\OLD\GL.DAT
    ACCT_3      SYS:DATABASE\OLD\AR.DAT
}

Customer Database            # Second set of database files
{
    CUST_1      SYS:CUSTOMER\SALES.DBF
}
```

In the above example, three database sets were defined. Each set definition has the following parts: (1) a set name; (2) an optional comment line; (3) an opening bracket; (4) a logical or alias file name; (5) an absolute path with file name; and (6) a closing bracket. These parts are illustrated below.



Backing up and Restoring Files

You back up Btrieve database sets just like you would any other file. Check the object detail area of the **Selection** tab of a backup job to see that the database set you want to back up is selected.

You restore database sets in the same way as other files, with one exception: you cannot restore a Btrieve database file with a different name without editing the TW\$BTRV.DAT file. (Note that you can't save the database *set* with a new name since the set only appears in TapeWare's database, not in the NetWare directory.)

To restore database sets to a *different* folder, on the **Selection** tab of a restore job, simply drag the database set to the new location. If you want to restore the database set in a *new* folder, use the **New Object** button to create a new folder and then drag the database set to that new folder.

If you need to restore a Btrieve database file with a different name, the easiest solution would be to first restore it to a new location and then change its name in the operating system. However you can change the name by editing the TW\$BTRV.DAT file. In this case, leave the logical or alias name unchanged, but modify the absolute path and file name. (Be certain to change the TW\$BTRV.DAT file back to its original form if you wish to continue to back up the original file.)

Additional Notes

You can specify any number of database sets. Each set can contain as many as 255 files.

The logical or alias file name is used by TapeWare to track a file, regardless of its physical location on a volume. Each logical name must be less than 48 characters long.

BTrieve files can be open and in use while TapeWare is backing up database sets. However, you must exit any BTrieve application before restoring a database set.

Configuring Email Support

TapeWare can be configured to automatically email the log of a job to that job's owner after the job has run. To do so, you must install one of the optional Email package and then enter a valid email address on the job owner's **Email** tab.

Note Install the Email support on the storage management server only.

Installing Email

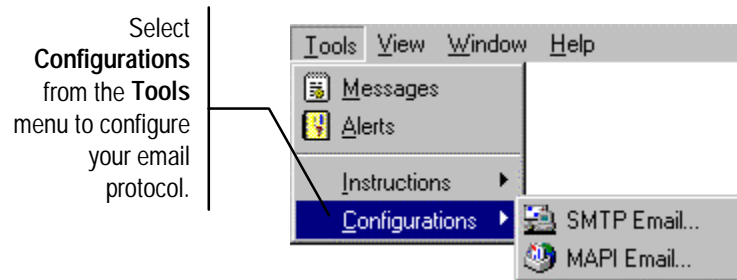
To install an optional email package, start the Installation Manager. Platform specific instructions are found in the section "Installation Instructions," Chapter 2.

In the Installation Manager window, select Install Package. The Installation Manager will guide you through the proper installation steps. You can install either the MAPI Email Package or the SMTP Email Package-or you can install both, although they must be installed during separate installation procedures.

Note If you install both email packages, duplicate emails of each job log will be sent.

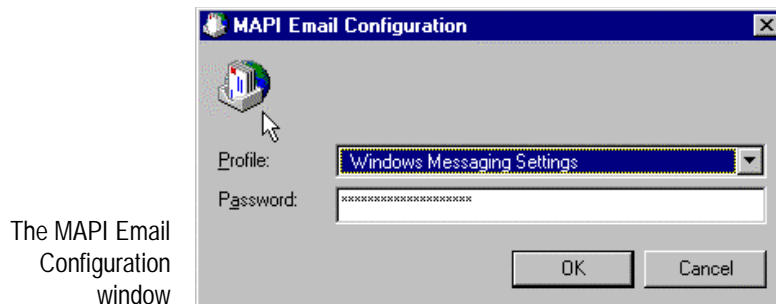
Configuring Email Packages

After installing an email package, you must configure the email. Select **Configurations** from the **Tools** menu and then either **MAPI Email** or **SMTP Email**, depending on your installation.



The MAPI Email Configuration Window

If your machine's operating system supports MAPI email protocol, you can use this protocol to email job logs.

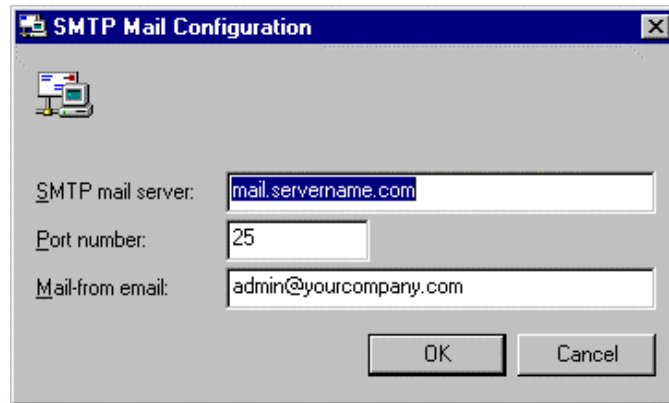


In the **Profile** field, select a profile. These profiles are set up independently by your operating system. If no profile is listed, your operating system has not been configured with an appropriate MAPI profile. Consult your email software documentation in order properly configure a MAPI profile. In the **Password** field, type in the appropriate password for the MAPI profile you have selected.

The SMTP Email Configuration Window

If your machine's operating system supports SMTP email protocol, you can use this protocol to email job logs.

The SMTP Email
Configuration window



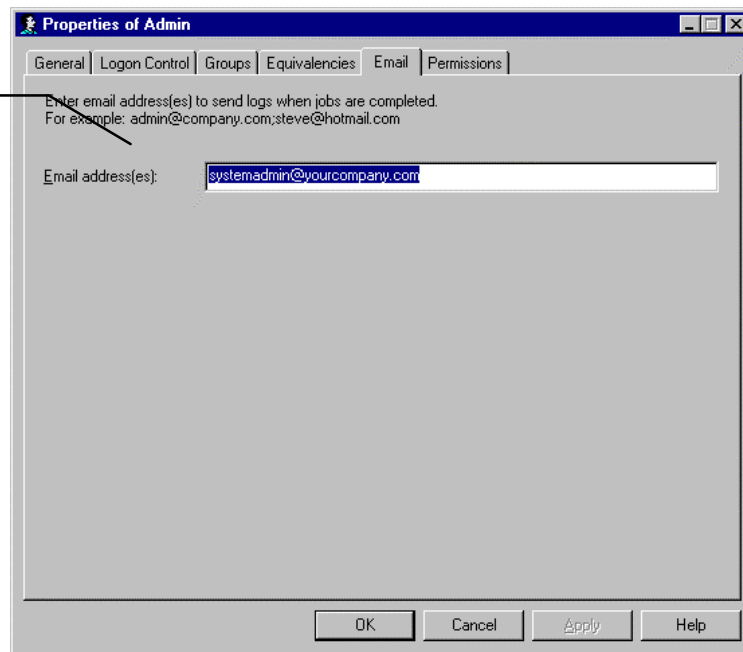
In the **SMTP mail server** field, enter the name of your mail server. In the **Port number** field, enter the appropriate SMTP port. By default, the **Port number** is set to 25. Generally, this default value is correct. If you are using a proxy server, however, you may have to enter a different **Port number**. In the **Mail from email** field, enter the email address that the SMTP server should send the email *from*. This must be a valid email address.

Note Some SMTP Email servers require that the Mail from be a valid user@host name, others ignore this field.

The Email Tab

When the email package is installed, a new tab, the **Email** tab, appears on the property sheet of each User object. TapeWare emails the job log to the address listed on the job owner's **Email** tab.

The job log is emailed to the address listed on the owner's **Email** tab.



The **Owner** of a job is generally the person who scheduled the job. You can confirm who owns the job by checking the **Owner** field on the **Queue** tab. If a scheduled job is forced to run, however, the User who forces the job to run becomes the job's temporary new owner. After this forced job has run, the owner of the job is reset to the last person who changed the job properties.

The job log is sent to whatever addresses are listed on the **Email** tab of the job's owner. You can specify multiple addresses on the **Email** tab; separate each address with a semicolon (no spaces). Additionally, you can enter the same address on every User's **Email** tab, and thus send a copy of *every* job log to the same user. For example, you might send the TapeWare administrator the log of every job that runs. Alternatively, consider setting up a separate email account for the sole purpose of receiving job logs.

Working with Microsoft Exchange Server databases

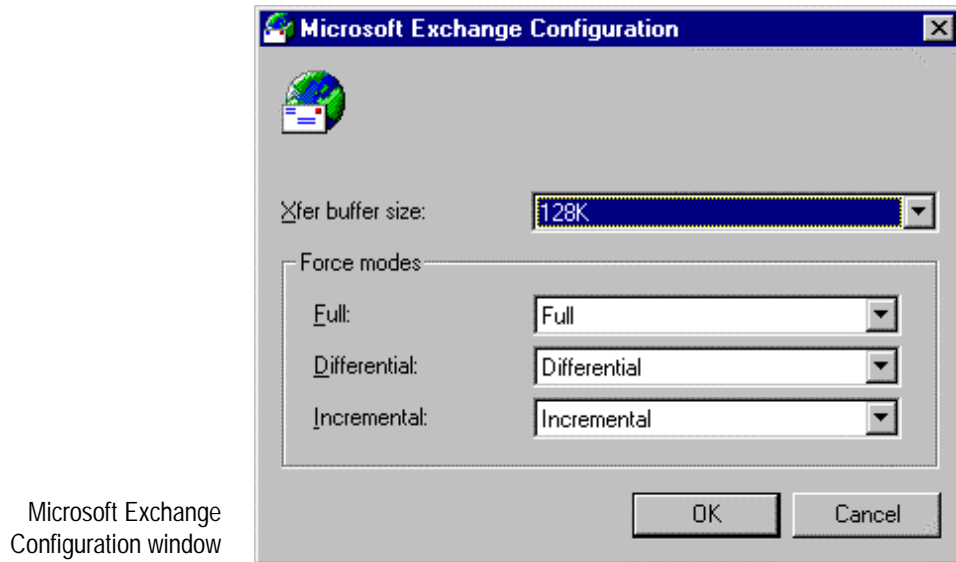
You can use TapeWare to backup Microsoft Exchange Server databases and configuration data. Acquaint yourself with the contents of the following chapter and create a working plan to deal with the disastrous loss of data now. TapeWare can help you minimize potential downtime and recover data quickly, but you must test and formulate a disaster recovery plan now.

In This Appendix

- Microsoft Exchange Configuration Window
- Microsoft Exchange Server Notes
- Restoring Microsoft Exchange Databases

Microsoft Exchange Configuration Window

Use the **Microsoft Exchange Configuration** window to set certain parameters which control how TapeWare works with Microsoft Exchange.



Transfer buffer size

This parameter sets the size of the read or write request TapeWare makes from Microsoft Exchange. The default value is **128K**. You can set the value anywhere from 64K to 1024K, in 64K increments.

The **Transfer buffer size** affects the speed at which data is transmitted and thus the speed at which jobs run. In general, you should leave the default value unchanged. In some installations, however, you may get better results by increasing the buffer size.

Force modes

As explained in the next section, the **Backup mode** parameter of a backup job affects Microsoft Exchange Server database files differently than file types. The **Force modes** parameters control how TapeWare backs up the database files.

Note that the settings here are only applicable to Microsoft Exchange Server database files; all other file types are backed up in the job's default mode. For example, if the **Backup mode** of a job is set to **Incremental**, and the **Force modes** parameter for incremental jobs is set to **Full**, TapeWare will back up the

Exchange Server database files in **Full** mode, but all other file types in **Incremental** mode.

Tip You can use this feature to insure that the database is always backed up in full mode, but that other files are only backed up when changed. This guarantees the greatest security for the most crucial files (that is, the Exchange Server database files), while not making jobs unnecessarily large by *not* backing up the entire network (that is, by backing up only the changed files).

Full: When the **Backup mode** of a job is set to **Full**, TapeWare checks this parameter to see how the job should be run with SQL database files. There is only one possible parameter, **Full**, and so the database files will be backed up in this mode. In this case, both the database and the transaction logs are backed up.

Differential: When the **Backup mode** of a job is set to **Differential**, TapeWare checks this parameter to see how the job should be run with Exchange database files. By default, TapeWare runs the jobs as a differential job, and so only the transaction logs are backed up.

If you want jobs with a **Differential** backup mode to backup both the database *and* the transaction logs, change this parameter to **Full**. In this case, TapeWare will treat the Exchange Server database files as if it were running a job in **Full** backup mode.

Incremental: When the **Backup mode** of a job is set to **Incremental**, TapeWare checks this parameter to see how the job should be run with Exchange database files. By default, TapeWare runs the jobs as an incremental job, and so only the transaction logs are backed up.

If you want jobs with an **Incremental** backup mode to backup both the database *and* the transaction logs, change this parameter to **Full**. In this case, TapeWare will treat the Exchange Server database files as if it were running a job in **Full** backup mode.

Microsoft Exchange Server Notes

When you use TapeWare to backup and restore Microsoft Exchange Server databases, you must pay special attention to the role Windows NT security serves in Microsoft Exchange and the backup mode of the TapeWare backup jobs.

Microsoft Exchange and Windows NT

Microsoft Exchange uses Windows NT security information for authentication, and thus when planning a comprehensive backup program, you must consider the Windows NT operating system as well. Be certain to include backup and restoration of the Windows NT operating system as part of your Microsoft Exchange disaster recovery plan.

Backup Modes

You can use the **Options** tab of a job to set the **Backup mode** to any of types of backup jobs: *full*, *incremental*, *differential*, or *[snapshot]*. For scheduled automatic rotation jobs, TapeWare automatically updates this job parameter to the value indicated on the **Schedule** tab of the job. For further information, see “Backup Options Automatically Updated” in Chapter 7 of the *User’s Guide and Technical Reference*.

When the **Backup mode** is set to **Full**, all files selected are backed up, including the entire information store and directory databases. Transaction logs are also backed up and then purged.

When the **Backup mode** is set to **Incremental**, only changes that have occurred since the last backup job are written. In particular, for database files, only the .log files are included in the backup job. *These .log files are then purged.*

When the **Backup mode** is set to **Differential**, for database files, only the .log files are included in the backup job, *but these files are not purged.*

When the **Backup mode** is set to **Snapshot**, TapeWare runs the job in **Full** backup mode. Note that this will cause the transaction logs to be reset (truncated). For this reason, running a job in **Snapshot** mode can compromise your comprehensive backup strategy if you are not careful to archive the media created by these jobs.

Backup Modes and Circular Logging

Microsoft Exchange Server supports database circular logging. Circular transaction logs differ from normal logs in that only a few log files are maintained. These files are purged automatically as new log files are created. When the transactions in the circular log files are recorded in the database, the log file is then deleted. New transactions are recorded in newly created log files.

If circular logging is enabled, *you cannot do incremental or differential backups.* These backup modes rely upon past transaction logs and thus are not available when circular logging is enabled. When circular logging is enabled, TapeWare will revert to *full* backup mode.

You can check to see if circular logging is enabled for a particular server by examining the **Advanced** tab of that server's **Properties** window. If you turn circular logging off, Microsoft Exchange Server will stop the database service and restart it after making the changes.

Restoring Microsoft Exchange Databases

To restore a Microsoft Exchange Server database, you must restore the database and all of the log files created since the last full backup job. To do so, you either (1) restore the database from the last full backup *if the last backup (the previous day's) was a full backup*; (2) restore the database from the most recent full backup and the last differential backup *if the last backup was a differential backup*; or (3) restore the database from the last full backup and all of the *incremental* backups made between that day and the present day.

Note that when you restore the database, you must create and run a separate job for each set of transaction logs you need to restore. You cannot skip any logs and the logs must be restored in sequential order. Thus, when recreating a database, you must first restore the whole database (created by a backup job running in *full* backup mode). Next, you must restore the transaction logs in the order created *and* in separate jobs. No log can be skipped when restoring.

For example, if you did a *full* backup on Monday, and *incremental* backups each day Tuesday through Friday, in order to restore the database to its state at the close of business Friday, you must run five separate jobs: one restoring the database from Monday's full backup job, and then four additional *separate* jobs restoring each transaction log in sequential order, beginning Tuesday and continuing with each log sequentially until Friday.

Microsoft Exchange Server Database Instances To check to see if a particular instance of a database is the whole (full) database or just the database log, open the **Instance** window of the database by selecting the database and clicking the **Instance** button on the toolbar of the **Selection** tab of the restore job. Highlight the instance you wish to investigate and click the **Details** button. In the **Stream format** field, it will show either **Database Full Stream** or **Database Log Stream**, depending on whether the database instance is the entire (full) database or only a transaction log.

◆ To restore a Microsoft Exchange Server Database

1. Find the date of the last full backup of the database.

To do so, select the database and click the **Check** button. In the **Instance** window, click the **Details** button. The **Stream format** field will show either **Database Full Stream** or **Database Log Stream**. Sequentially move

through the instances in the **Available instances** field by date until you find the most recent full backup of the database.

That instance will be selected for restoring when it is highlighted in the **Available instances** field. Click **Ok** to restore that instance.

2. Create and run a restore job of the most recent instance of a full backup of the database.
3. If the most backup was a full backup, skip the rest of these steps and restart the Microsoft Exchange Server database. As the service is restarted, it automatically restores all of the transactions from the transaction logs.
4. If the most recent backup job was a differential job *and you have performed no incremental jobs between the date of the last full backup and the most recent backup*, then create and run a new restore job, selecting the **<Latest>** instance of the database. Then restart the Microsoft Exchange Server database. As the service is restarted, it automatically restores all of the transaction from the transaction logs.

Note: if you have performed any incremental jobs since the date of the last full backup, follow the instructions in the next step (5).

5. If you have run an incremental backup job after the most recent full backup job, you must create and run a separate restore job for each backup performed after the most recent full backup. Sequentially select instances of the database from the **Available instances** field in the **Instances** window of the database. Run and complete each restore job before creating and running a new restore job.

Continue to create and run restore jobs until you have restored the **<Latest>** instance of the database. Then restart the Microsoft Exchange Server database. As the service is restarted, it automatically restores all of the transaction from the transaction logs.

Working with Microsoft SQL Server Databases

This Appendix contains important information pertaining to backing up and restoring Microsoft SQL Server databases. If you are using TapeWare to backup and restore SQL Server databases, be certain to carefully read and follow these instructions.

In This Appendix

- Overview
- SQL Configuration Window
- SQL Server Backup Job Notes
- SQL Server Restore Job Notes
- Restoring SQL Server User Databases
- Restoring SQL Server Master Databases

Overview

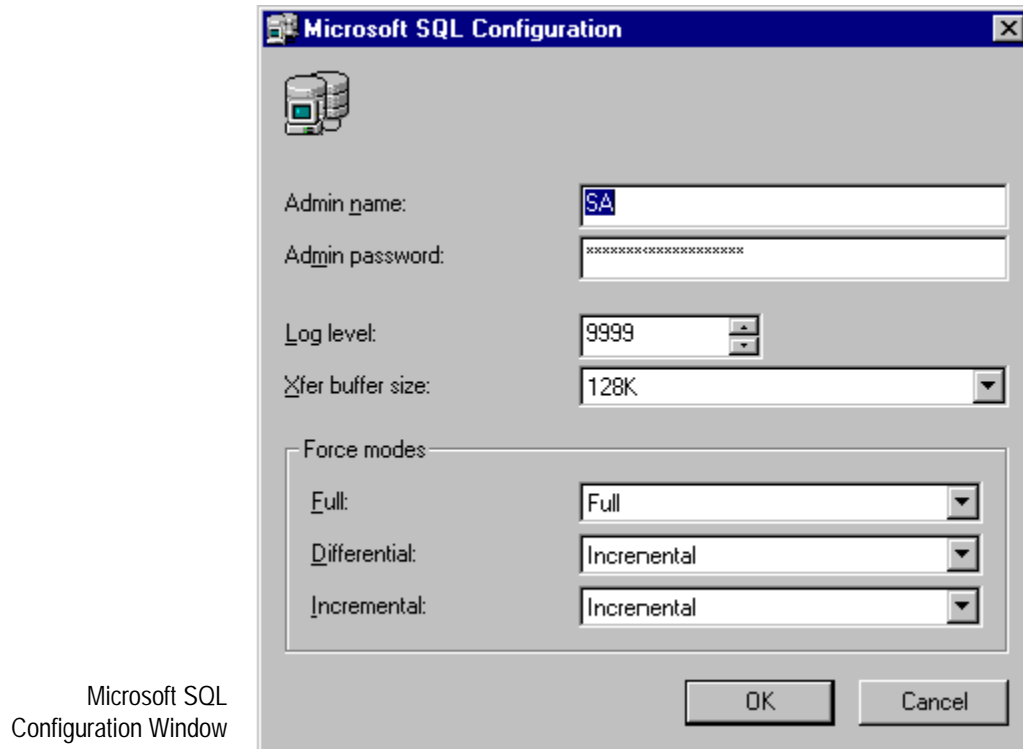
Many SQL server environments are mission-critical and must be maintained twenty-four hours a day, seven days a week. Procedures and plans must be in place to insure the quick recovery of data in the event of data loss.

Key to the quick recovery of databases are the transaction logs. By using the transaction logs associated with each database, quick recovery of a database can be achieved. Transactions that were not committed can be rolled back, while transactions that were committed can be written to disk.

While transaction logs assure that only committed transactions are written and restored, in order to use them correctly, you must have a comprehensive backup plan that regularly backs up these logs. Additionally, when you reconstruct a database, you must restore the database and logs using only the procedures set out below.

SQL Configuration Window

Use the **Microsoft SQL Configuration** window to set certain parameters which control how TapeWare works with SQL Server.



Admin name

TapeWare will send this name to Microsoft SQL Server whenever an user name is required. Type the name of the Microsoft SQL administrator in this field. The default name is **SA**.

Admin password

TapeWare will send this name to Microsoft SQL Server along with the administrator's name whenever required. There is no default value.

Log level

This parameter controls which error messages are sent by Microsoft SQL. The lower the number is set, the lower the severity of the error message required before the message is sent. Thus lowering the number increases the frequency of messages from Microsoft SQL.

The messages sent to TapeWare can be viewed in the **Message** window, available from the **Tools** menu. Alternatively, these messages are stored in the log of a TapeWare job. The log of any job can be viewed from the **Logs** tab of that job.

The default value is 9999. In general, you should leave this value unchanged.

Xfer buffer size

This parameter sets the size of the read or write request TapeWare makes from SQL. The default value is **128K**. You can set the value anywhere from 64K to 1024K, in 64K increments.

The **Xfer buffer size** affects the speed at which data is transmitted and thus the speed at which jobs run. In general, you should leave the default value unchanged. In some installations, however, you may get better results by increasing the buffer size.

Force modes

As explained in the next section, the **Backup mode** parameter of a backup job affects Microsoft SQL Server database files differently than file types. The **Force modes** parameters control how TapeWare backs up the database files.

Note that the settings here are only applicable to Microsoft SQL Server database files; all other file types are backed up in the job's default mode. For example, if the **Backup mode** of a job is set to **Incremental**, and the **Force modes** parameter for incremental jobs is set to **Full**, TapeWare will back up the SQL Server database files in **Full** mode, but all other file types in **Incremental** mode.

Tip You can use this feature to insure that the database is always backed up in full mode, but that other files are only backed up when changed. This guarantees the greatest security for the most crucial files (that is, the SQL Server database files), while not making jobs unnecessarily large by *not* backing up the entire network (that is, by backing up only the changed files).

Full When the **Backup mode** of a job is set to **Full**, TapeWare checks this parameter to see how the job should be run with SQL database files. There is only one possible parameter, **Full**, and so the database files will be backed up in this mode. In this case, both the database and the transaction logs are backed up.

Differential When the **Backup mode** of a job is set to **Differential**, TapeWare checks this parameter to see how the job should be run with SQL database files. By default, TapeWare runs the jobs as an incremental job, and so only the transaction logs are backed up. *There is no distinct **Differential** job mode for SQL Server databases.*

If you want jobs with a **Differential** backup mode to backup both the database and the transaction logs, change this parameter to **Full**. In this case, TapeWare will treat the SQL Server database files as if it were running a job in **Full** backup mode.

Incremental When the **Backup mode** of a job is set to **Incremental**, TapeWare checks this parameter to see how the job should be run with SQL database files. By default, TapeWare runs the jobs as an incremental job, and so only the transaction logs are backed up.

If you want jobs with an **Incremental** backup mode to backup both the database and the transaction logs, change this parameter to **Full**. In this case, TapeWare will treat the SQL Server database files as if it were running a job in **Full** backup mode.

SQL Server Backup Job Notes

Two additional concerns are present when backing up SQL Server databases: setting the **Backup mode** of a job to either **Full**, **Incremental**, or **Differential**; and configuring TapeWare to work with SQL Server's built-in backup routine.

SQL Server Databases and the Backup Mode

The **Backup mode** on the **Options** tab of job that backs up SQL Server databases is especially critical and important.

Backup Modes

When the **Full** parameter is selected, all files selected for backup are backed up, including SQL Server databases. However, when either the **Incremental** or **Differential** option is selected, TapeWare backs up only the transaction logs. *There is no difference between **Incremental** and **Differential** jobs for SQL Server databases.*

When the **Backup mode** is set to **Snapshot**, TapeWare runs the job in **Full** backup mode. Note that this will cause the transaction logs to be reset (truncated). For this reason, running a job in **Snapshot** mode can compromise your comprehensive backup strategy if you are not careful to archive the media created by these jobs.

Additional Conditions

The **Backup mode** is subject to the following additional provisions:

- Master, Model, MSDB, and Pubs databases support only full backups. The **Backup mode** option is automatically set to **Full** when backing up these databases.
- If the Log Device is the same as the Database Device, only full backups are allowed. The job will always run as a full backup, even if the **Backup mode** is set to **Incremental** or **Differential**.

If you set a job to run in either **Incremental** or **Differential** mode and the job can only run as a full backup (as a result of one of the provisions above), the job will fail to run on each of its initial passes, but will run in **Full** backup mode on its final pass.

Using TapeWare with SQL Server's Backup Routine

SQL Server has built-in utilities and commands for backing up data. When you use TapeWare to backup SQL Server databases, you can still use these built-in SQL Server utilities and commands.

For example, you can use the DUMP command to dump transaction logs to the dump device (preferably, a separate disk drive). You can set this up to occur at regular intervals, such as every 15 minutes or every hour. Next, you can create a backup job that backs up these transaction logs onto archival media every day.

In general, when you implement TapeWare to back up your SQL Server databases, continue to use SQL Server's internal commands to duplicate and backup transaction logs. Set up a separate TapeWare backup job to write these duplicated transaction logs to archival media.

SQL Server Restore Job Notes

When restoring SQL Server databases, you must (1) restore the logs in the order created, (2) restore databases to the appropriate original device, and (3) follow special procedures when renaming databases.

Restoring SQL Server Transaction Logs

When recreating a database, you must first restore the whole database (created by a backup job running in *full* backup mode). Next, you must restore the transaction logs in the order created *and* in separate jobs. No log can be skipped when restoring.

For example, if you did a *full* backup on Monday, and *incremental* backups each day Tuesday through Friday, you must run five separate jobs: one restoring the database from Monday's full backup job, and then four additional *separate* jobs

restoring each transaction log in sequential order, beginning Tuesday and continuing with each log sequentially until Friday.

You do not have to follow these procedures when restoring databases backed up with *full* backup jobs. (**Full** backup jobs back up the entire database, while **Incremental** and **Differential** jobs only back up the database logs.)

SQL Server Database Instances To check to see if a particular instance of a database is the whole (full) database or just the database log, open the **Instance** window of the database by selecting the database and clicking the **Check** button on the toolbar of the **Selection** tab of the restore job. Highlight the instance you wish to investigate and click the **Details** button. In the **Stream format** field, it will show either **Database Full Stream** or **Database Log Stream**, depending on whether the database instance is the entire (full) database or only a database log.

Restoring SQL Databases to Devices

When restoring databases, if TapeWare discovers that the database already exists, it restores the database to the current device(s) (that is, in the appropriate disk partition, etc.). *New* devices are created only when the original database is no longer available.

To accomplish this, TapeWare first determines if the database exists. If the database does exist, TapeWare will use that database. (In this case, TapeWare doesn't check to see if database device is the original device.)

If the database does *not* exist, TapeWare next identifies the *database devices* where the database was originally located. If these database *devices* already exist, TapeWare will restore the database to those devices.

If a database *device* does *not* exist, TapeWare *recreates* that database device at its *original* location and with its original size. After all the database devices are created, TapeWare then creates the database with all the original options at the original locations.

As a result, if the original device is no longer available, it will be recreated. TapeWare, however, will only recreate the device if the same physical disk drive is available (that is, in a physical device with the same designated drive name, and so forth). Thus, if the drive is not available for some reason (for example, it is off-line or is corrupted), the restore job will fail.

Suppose, however, that you wish to restore the database to a new location. For example, you originally had the database on two separate 1-gigabyte drives and now want to locate the database on a new 5-gigabyte drive. To do this, use SQL Server to set up the database and database device in the new location, then run a TapeWare restore job. TapeWare will find that the database exists and will restore it to that device, even though the database is in a new location.

Restoring SQL Databases with a New Name

You can rename a database while restoring using the normal procedures for renaming files outlined in the *User's Guide and Technical Reference*. This method involves changing the name on the **Selection** tab of the restore job. However, *you cannot rename the master database*. When you restore a master database, you must follow the procedures specified in the last section of this Appendix.

Restoring SQL Server User Databases

To restore a database, begin by restoring the most recent **full** backup of the database, followed by *all* the database logs, that is, backups made with the **Backup mode** set to either **incremental** or **differential**.

To check to see if a particular instance of a database is the whole database or just the database log, open the **Instance** window of the database by selecting it and clicking the **Check** button on the toolbar of the **Selection** tab of the restore job. Highlight the instance you wish to investigate and click the **Details** button. In the **Stream format** field, it will show either **Database Full Stream** or **Database Log Stream**, depending on whether the database instance is the entire (full) database or only a database log.

Note When the stream format type is “Database Full Stream,” this is equivalent to the SQL statement DUMP DATABASE. When the stream format type is “Database Log Stream,” it is equivalent to the DUMP TRANSACTION statement.

When a database is restored, if the database does not yet exist, TapeWare will create the database on the devices where the database was originally located. If these database devices do not exist, TapeWare will automatically create the database devices required for each database before creating the database.

Note If the database already exists, make sure the database devices are all valid and ready. If some of the database devices are not found, or if your database is in the “suspect” state, drop the database and any database devices which are having problems. These databases and devices will be recreated by TapeWare when the database is restored.

◆ To Restore a Lost or Damaged Database

1. If the transaction log of the damaged or inaccessible user database is on an undamaged device, make a backup of the transactions before proceeding. (This allows you to preserve up to the minute transactions that are not included on the backup tape.)

You may use either a DUMP TRANSACTION statement on the SQL server, or use a TapeWare **Incremental** backup job to back up only the transactions logs.

2. If you are restoring the database because the data in the database is no longer needed or is incorrect, skip to the next step (step 3). The following instructions are for recreating database devices and the database which had existed previously.

During the restore processes, TapeWare will recreate the database and all segments exactly as they existed when the backup was performed.

To do this, TapeWare first determines if the database exists. If the database does exist, TapeWare will use the database as is *without any further processing or changes*.

If the database does not exist, TapeWare next identifies the database devices on which the database was originally located. If the appropriate database *device* already exists, TapeWare will use that device as is without further processing.

If the database *device* does not exist, TapeWare *recreates* the database device at its *original* location and with its original size. After all the database devices are created, TapeWare then creates the database with all the original options at the original locations.

Tip This method make disaster recovery simple. The user need merely create a restore job and allow TapeWare to recreate whatever is needed in order to successfully restore the database.

Note, however, if a disk drive fails and is not replaced, TapeWare will be unable to restore your database because it will be unable to recreate a database device. For example, if a segment of your database resides on a database named 'DATA' at D:\MSSQL\DATA\DATA.DAT, if D: is lost and not replaced, when TapeWare attempts to recreate the database device, it will fail, since D: no longer exists.

To avoid this problem, manually recreate the database device at some other location. It must be at least as large as the original database device since

TapeWare will attempt to create a database segment on it the same size as the original database.

An alternative method is to manually create the entire database itself. Thus, when TapeWare attempts to restore the database, since the database already exists, it will use that preexisting database. This allows you to restore a database in a new location, since TapeWare does not check to see if it is the original device before restoring the database, because the database already exists.

Note TapeWare tracks databases and database devices *by name*. So, if a database or database device already exists with the same name, TapeWare will use that database or device.

3. Using TapeWare, create a restore job and run the job to restore the database. You must start with an instance of the database to restore which was created using a **Full** backup job. Check the **Stream format** field in the **Instance** window of the database to verify that it was a **Full Database Stream**.
4. Create additional restore jobs to restore each transaction log backed up after the full database you restored. You must create and run a *separate restore* job for each transaction log.

For example, if you ran a full backup on Friday, and incremental jobs (that is, jobs that backed up only the transaction logs) on the following Monday and Tuesday, you must first restore the database using Friday's instance of the database. Next, create a run and restore job that restores Monday's instance (Monday's transaction log). Finally, create and run a job that restores Tuesday's instance (Tuesday's transaction log).

Restoring SQL Server Master Databases

A damaged master database is evident by the failure of the SQL Server to start, by segmentation faults or input/output errors, or by a report from DBCC. An example of an error might be damage caused by media failure in the area in which master database is stored.

The procedure used to recover a damaged master database is different from the procedure used to recover user databases. If the master database becomes unusable, it must be restored from a previous dump. All changes made to the master database after the last dump are lost when the dump is reloaded and therefore must be reapplied.

It is strongly recommended that the master database be backed up each time it is changed. This is best accomplished by prohibiting the creation of user-defined objects in the master database and by being aware of the statements and system procedures, and the equivalent actions in SQL Enterprise Manager, that modify it.

The most common statements and system procedures that modify master are:

- DISK INIT
- CREATE DATABASE
- ALTER DATABASE
- DISK MIRROR
- DISK UNMIRROR
- DISK REMIRROR
- sp_dropremotelogin
- sp_addumpdevice
- sp_dropdevice
- sp_addlogin
- sp_droplogin
- sp_addserver
- sp_dropserver
- sp_addremotelogin

If a user database is created, expanded, or shrunk after the most recent dump (backup) of the master database, and if it becomes necessary to reload the master database, then that user database and all data in will be lost and must be restored from backup. Because of this, *always dump (backup) the master database after creating, expanding, or shrinking user databases.*

◆ To Recover a Damaged Master Database

1. Use the SQL Setup program to rebuild the master database.

You must rebuild using the same character set and sort order as the master database dump that will be reloaded.

2. Restart SQL Server in single-user mode.
3. Restore the master database from the most recent backup.
4. Apply to the master database any changes that were not included in the most recent backup.
5. Restore the msdb database.

These steps are explained in the sections that follow.

Step 1 - Rebuild the master Database

1. From the Microsoft SQL Server program group, double-click the **SQL Setup** icon.

(Alternatively, from the distribution media, from the directory containing the software compatible with your hardware platform's processor architecture, run SETUP.EXE.)

2. Respond to the on-screen instructions until the **Options** dialog box appears.
3. Select the **Rebuild Master Database** option, and then choose the **Continue** button. A confirmation dialog box will appear.
4. Choose **Resume**. The **Rebuild Options** dialog box appears.
5. To specify the character set, choose the **Sets** button and complete the **Select Character Set** dialog box that appears. (If you will be using the default character set (ISO 8859-1), skip this step.)

Note You must use the same character set and sort order that were previously used for this master database.

6. To specify the sort order, choose the **Orders** button and complete the **Select Sort Order** dialog box that appears. (If you will be using the default sort order (dictionary order, case-insensitive), skip this step.)
7. In the **Rebuild Options** dialog box, choose **Continue**. The **SQL Server Installation Path** dialog box appears.
8. If not correctly displayed in the **SQL Server Installation Path** dialog box, enter the location of the existing SQL Server installation, and then choose **Continue**.

The **Rebuild Master Device** dialog box appears.

9. If it is not correctly displayed in the **Rebuild Master Device** dialog box, enter the location and name of the existing MASTER device. Also enter a MASTER device size, and then choose Continue.

The setup program will then rebuild the master database.

10. When rebuilding is complete and the completion dialog box appears, choose the Exit button.

Note The files MASTER.DA@ and MASTER.AL@ are stored in the \MSSQL\INSTALL directory. When rebuilding the master database (or when installing SQL Server), one of these two files is used by the setup program. When the default sort order and character set are selected, MASTER.DA@ is expanded and copied onto the server, replacing MASTER.DAT. When an alternate character set and/or sort order is selected, MASTER.AL@ is expanded, copied onto the server, and several SQL scripts are run.

Step 2 - Restart SQL Server in Single-user Mode

Before you can restore the master database, you must start SQL Server in single-user mode.

1. If SQL Server is already running, stop the server. You can perform this operation from either the SQL Enterprise Manager or from the NT Service control panel.
2. From a command prompt, type:

```
SQLSEVR /c /dmaster_device /m
```

where

```
/c
```

starts SQL Server independent of the Windows NT Service Control Manager,

```
/dmaster_device_ path
```

specifies a physical name for the MASTER database device, and

```
/m
```

specifies single-user mode. Consider this example:

```
C:\MSSQL\BIN> SQLSERVER /c /dC:\MSSQL\DATA\MASTER.DAT  
/m
```

Note SQLSERVER.EXE is usually located in \MSSQL\BIN.

Step 3 - Restore the Master Database from the Most Recent Backup

1. Create a restore job and select the most recent instance of the master database.
2. Run the restore job.

Note This may take some time, typically 10–15 minutes depending on the size of the master database. Restore only the master database while in single user mode. Do not restore any other databases.

Step 4 - Apply Changes to the Master Database

1. Restart the SQL server if it is not already started. You can perform this operation from either SQL Enterprise manager or the NT service control panel.

If there have been no changes to the master database since the last dump, then proceed to step 5.

2. If login IDs or devices have been added to or dropped from the master database since the last backup, those changes must be reapplied. Restart the server and reapply the changes manually or from saved batch files.
3. If databases have been created, expanded, or shrunk since the last dump of master, those databases must be dropped and then restored.
4. If you have made many changes and have no recent dump, it is possible that by reloading master in some cases you can regain data in user databases that has been lost. This technique requires the use of DISK REINIT and DISK REFIT and can involve manual modifications to the master database tables.
 - Use DISK REINIT to re-create rows in sysdevices for all database devices that have been added after the most recent dump. DISK REINIT updates sysdevices just as DISK INIT does, but it does not format the physical disk file, so existing data is preserved.
 - Use DISK REFIT to re-create rows in sysusages and sysdatabases for all CREATE and ALTER DATABASE statements that were performed after the most recent dump.

DISK REFIT scans the physical file associated with each space that is allocated to databases. It also adds the corresponding sysdatabases entries. Some of the information is not reconstructed perfectly. For example, the original virtual device number is not assigned, because it is not known. Instead, virtual device numbers are assigned sequentially. The database owner is not extracted while scanning the physical files; ownership is set to the system administrator. It is also not possible to determine how many sysusages entries originally existed. DISK REFIT inserts a separate entry for each different segment type.

- When this is done, correct the entries made by DISK REFIT to sysdatabases and sysusages (if desired) and also add to syslogins any login IDs that were not retained. Then shut down and restart SQL Server.

Warning Capturing the latest changes made to a database by using DISK REFIT and DISK REINIT to re-create the master database is possible, but it is preferable to keep the master database current by dumping it after creating or altering databases. Using DISK REFIT and DISK REINIT is a complicated process that can result in data loss because many of the changes made to a database often must be reconstructed manually in the master database. If you feel this technique is necessary, contact your primary support provider before beginning the recovery process.

Step 5 – Drop Invalid Databases And Database Devices

1. Use the SQL Enterprise manager to drop any invalid database devices and databases from the newly restored master database.

Note If you are recovering from a disaster where you have lost a database device file, the master database you have just restored still contains a reference to it. TapeWare will not be able to restore any databases contained on the database device until the file is restored or the database device is dropped. If the database device is dropped, TapeWare will automatically recreate the device when a database contained on the device is restored.

Step 6 - Restore the msdb Database

Procedures for restoring a SQL Server databases were covered in the previous section.

When restoring a msdb database, keep the following considerations in mind:

-
- The msdb database supports SQL Executive and provides a storage area for scheduling information. The schedules that you implement using SQL Enterprise Manager are maintained in the msdb database. This includes such things as the tasks that you schedule from the Task Scheduling window, the automatic backups you schedule from the Database Backup/Restore window, and all replication tasks (which are automatically created by the system if the server is configured as a replication distributor).
 - During installation of a server, the setup program automatically creates two devices (of 2 MB and 1 MB) on the same disk drive as the master database, and then places the msdb database on the 2 MB device (MSDBDATA) and its transaction log on the 1 MB device (MSDBLOG). Scheduling information is then stored in this database.
 - During a rebuild of the master database, the setup program drops and re-creates the msdb database, which results in a loss of all scheduling information.

Troubleshooting Guide

This chapter contains useful information about commonly encountered problems and frequently asked questions about implementing TapeWare.

In This Appendix

- Troubleshooting Installation
- Troubleshooting Backup Jobs
- Troubleshooting Restore Jobs
- Troubleshooting Verifying
- Troubleshooting Backup Devices
- Troubleshooting the Storage Management Database
- Troubleshooting Error Messages

Troubleshooting Installation

1. I entered my Key Code correctly, but the installation won't continue.

Make certain that the Key Code you entered supports the platform on which you are trying to install TapeWare.

Each Key Code is product and configuration specific. Some Key Codes allow you to back up unlimited machines and servers; others only allow you to back up a single desktop machine. Other Key Codes work for both NetWare and Windows NT networks; others, for only one of the two.

Your license agreement determines the number of machines and servers on which you can install TapeWare and the operating system or platform of each workstation or file server. If you need to install TapeWare on additional machines, you can purchase an upgrade to your license agreement that will allow you to do so.

2. I selected 'Connect to an existing storage management server' during installation, but now that storage management server is not displayed.

First, make sure the storage management server you are attempting to connect to is running. Also, make sure that it has the necessary networking software that is operating properly. You can verify this by using the **Database** tab. Open the Network folder on that tab, and select the machine you intended to be the storage management server. If that machine has been configured for network operation, you will see a driver named Xpt-... in the object detail area of the tab (on the right side). If you don't see that driver, you must reinstall TapeWare on the storage server in order to add networking support to it.

Second, if you are running TCP/IP, TapeWare will only display servers on your local subnet. If you are not on the same subnet as the storage server, you must specifically add the address or host name of the server during the installation process.

3. When attempting to log on to NDS, TapeWare keeps returning to the password prompt.

You are not logging on to TapeWare at this point. You are actually logging on to the NDS tree. Therefore, you must enter the NDS password.

Additionally, TapeWare will attempt to display the default NDS information; sometimes, however, this information may be incorrect. If NDS logon displays something like "cn=admin.ou=organizationalUnit.o=organization," change the path to "cn=admin.o=organization."

4. I can't log on to TapeWare.

If this is a new TapeWare installation, first attempt to log on as the ADMIN user. This account is automatically created during installation. There is no password initially for the ADMIN.

If an 'error 25 - Service not available' error is displayed, you may not have installed TapeWare correctly. TapeWare requires a storage management server to be active in order for you to log on. During installation, make sure you selected "Create a new storage manager" on one of the machines you are installing, then start TapeWare on that machine.

Troubleshooting Backup Jobs

5. When I run a backup job, the backup device used is not the one on the server, but rather the one on my local machine.

On the **Options** tab of the job, you have probably accepted the default **Network** device selection. When this is enabled, TapeWare will use any device on the network, in this case, your local backup device.

If you want to send the job to a specific device, first delete the **Network** object in the **Device** field. Then click on the **Add...** button and add the specific device you want the job to use.

6. Does TapeWare support media rotation?

Media rotation is supported via a scheduling option. TapeWare creates rotation media when you select one of the pre-defined rotation schedules on the **Schedule** tab of backup job.

To create a rotation job, begin by selecting the files you want to back up. Then select one of the pre-defined rotation schedules on the **Schedule** tab. TapeWare will automatically create the rotation media and run the proper incremental and full backups for the schedule type you have selected.

For additional information, see "Which Built-in Schedule to Select," Chapter 6.

7. TapeWare is running, but I don't see any network mapped drives.

Under normal operating conditions, TapeWare requires that a copy of TapeWare be installed on all machines that are backup targets. When TapeWare is installed on a workstation or file server, it has access to all that machine's resources, such as the registry, bindery, NDS, and so forth. TapeWare needs these resources in order to properly back up all that workstation's or file server's data.

It is possible, however, to enable alternative network support. To do so, modify the “driveEnable=...” setting in the TapeWare.Ini configuration file. This method is not recommended, however. When backing up files using this method, critical information may not be backed up, such as trustee data, access control lists, owners, registries, and so forth.

8. My scheduled job is not running.

First, make sure the job is scheduled. Open the properties sheet of the job and check the **Schedule** tab. Be certain to close the properties sheet after examining the schedule; the job will not run when its properties sheet is open. Next, check the **Queue** tab to verify that the job is scheduled.

Next, if you are running under Windows and have not installed TapeWare as a service, you must keep TapeWare open in order for the job to run. TapeWare is shut down when you exit from the application and so must keep TapeWare running in order for the job to be executed at the specified time.

For additional information, see “Automatically Running Scheduled Jobs” in Chapter 8 and “Logging Out and Running Scheduled Jobs” in Chapter 2.

9. I have two tape drives, but only one is used to run a job.

TapeWare uses “streams” to divide a job up and assign it to devices for backup. By default, a new stream is created for each disk volume (such as, C:) to be backed up. *To use multiple devices, you must have multiple streams.* The streams are then evenly distributed across all available backup devices. If you have only a single volume, only one stream will be created by default.

To create additional streams, change the **Backup stream** parameter on the **Storage** tab of a directory to **Create new stream**. For further information, see “Storage Tab,” Chapter 12, and “Strategies for Faster Jobs,” Chapter 10.

One word of caution: since multiple streams run concurrently, creating multiple streams on the same physical disk drive does not necessarily result in faster backup jobs, since the drive will be attempting to stream to multiple devices at once, which will require numerous seek and read commands from various sectors at the same time.

10. How do I replace media in a rotation group?

TapeWare automatically creates a series of folders and media for use with rotation jobs. These folders control the daily, weekly, monthly and yearly tapes. To remove a tape if it is physically damaged, or because it is lost, delete the media from the set. TapeWare will automatically format replacement media when required.

To move media to an off-site location, begin by creating a new media folder in your User/Group folder. Give it a name like “Off-site Media”; then, drag the media to this folder. If the tape is required by TapeWare during the rotation schedule, it will automatically create a new tape to replace the tape you moved to off-site storage.

11. How can I tell when the next job will run and which media is required?

Click on the **Wizards** tab and select **Instructions**.

12. How can I easily view the logs for each Job?

Click on the **Wizards** tab and select **Instructions**. Tab to the **Logs** screen and view any available log.

13. How can I determine which files were not backed up?

From the **Logs** tab of the backup job, select the date of the backup you want to check. TapeWare uses Notepad or Wordpad to view the job log, allowing you to save the log to a file after editing. (Note that you are only working with a copy of the log. The original log is still available.) TapeWare can also print the log directly to the printer.

14. Can I print the reports or logs under NetWare?

You can print to any local print queue under NetWare. The queue must be local, but the printer need not be local. By configuring a remote print server to service the local queue, the print job can be routed to a remote printer.

Troubleshooting Restore Jobs

15. I can't restore a backup made from Windows NT to Windows 95/98.

Windows NT, Windows 95/98, DOS, and NetWare store information in different formats. For example, NetWare stores information on tape in a compressed format if the file is stored on disk in compressed format, but Windows cannot read NetWare compressed data.

In order to restore to a different operating systems, and sometimes, even to a different version of the *same* operating system, you must create the backup in a “compatible” format. To do so, open the **Advanced options** window of the job from the job's **Options** tab. The clear (turn off) the **Native data streams format** option. This causes TapeWare to decompress the data before it is backed up; however, you may lose security information under Windows NT.

For additional information, see “Moving Data between Operating Systems,” Chapter 10.

16. How can I restore data to a different file name?

Select the file, directory, or volume you wish to restore with a different name. Then open the properties sheet for that object. On the **General** tab, type in the new name and press **Ok**.

For additional information, see “Restoring Files with New Names,” Chapter 10.

17. How can I restore data to another location?

To restore data to another location, open the **Selection** tab of the restore job. Click on the directory or file you want restored to a different location, then drag that object to the new target location. To use the keyboard, use CTRL-C to copy the object and CTRL-V to move the object to its new location.

If you have not backed up the target directory (the place where you want to move the files to), it will not be displayed. In this case, click the right mouse button on the files or directories you want to move, and then select **Move...** from the shortcut menu. In the **Confirm Move** window, click on the **Browse...** button to select a new target location.

18. How can I restore all the files from a single session?

Every instance of a file or directory backed up during a single job has the same instance date. You can use this information to select all the files from a single session.

On the **Selection** tab of the restore job, select the file, directory or volume to restore. Press the **Select Instance** button on the tool bar to open the **Instances...** window. Select the date of the instance you want restored. All children (files) having the same instance will also be selected.

For additional information, see “Selecting Instances from a Specific Job,” Chapter 10.

19. How can I determine which files are on particular media?

Click on the **Wizard** tab, and select **Media Content**. Follow the tree and pick the media you want information about.

20. When restoring, I get many alerts. What is wrong?

If you select many devices or set the **Device** field on the **Options** tab of the job to **Network** (the default setting), TapeWare will attempt to complete the restore job using *all* the devices listed or found. If a device does not contain media, or the device contains the wrong media, an alert is sent. This alert informs you to put the desired media into the appropriate device.

For example, suppose you are restoring from a single media but the **Device** field on the **Options** tab lists 4 devices. If the target media is in the third device listed, TapeWare will issue 2 alerts, one for each of the first two devices. These alerts can be ignored.

Note as well that you do not have to insert the *requested* media into the device. To fulfill the request, any of the media the job requires can be placed in the device. TapeWare will then use the media that you actually did put into the device.

If you want the job to use a specific device, first delete the **Network** object in the **Device** field. Then click on the **Add...** button and add the specific device you want the restore job to use.

21. Does TapeWare back up files as compressed?

TapeWare will copy files to tape in compressed format without decompressing them first. This significantly enhances the speed of the backup.

22. Can I restore NT files to NetWare?

Yes, but you must clear the **Native data streams format** option from the **Advanced Options** window on the *backup job's Options* tab. When this option is cleared, the data will be written to the media in a generic format that NetWare can understand.

Note that the data must be *backed up* in this generic format; the **Native data streams format** option has no function for restore jobs.

For additional information, see “Moving Data between Operating Systems,” Chapter 10.

23. Can I restore NetWare files to NT?

Yes, but you must follow the same procedures specified above. Clear the **Native data streams format** option for the *backup job*. This causes NetWare data to be decompressed before being backed up; it can then be understood by NT, which does not understand NetWare compression.

Troubleshooting Verifying

24. I occasionally get a ‘stream sync error’ when verifying media.

This is usually caused by a physical read problem from the backup device. The data TapeWare expected from the media was not found. This can be caused by:

- *Bad media*: try replacing the media;

- *Read errors on the drive:* try cleaning the drive heads;
- *SCSI errors:* try checking the SCSI termination; or
- *Driver errors:* try checking to see that you are not using a real mode ASPI drive. Check your config.sys file for something like device=ASPI4DOS.SYS, ASPI8DOS.SYS, and so forth.

Troubleshooting Backup Devices

25. I can't see my tape drive on the Device tab.

TapeWare will automatically recognize any backup device attached to your file server or workstation, provided they are connect to the machine via a SCSI connection. If your machine recognizes the backup device as a valid SCSI device, then it should automatically appear as a backup device in the storage management database.

If the device you were expected does not appear, first check to see that your machine or network recognizes that device. Use the Windows Explorer (or similar program) and check to see that the backup device appears as a SCSI device on the machine.

Next, check to see if any device drivers had failures during loading. You can check this by selecting **Messages** from the **Tools** menu (or by pressing CTRL-F8 on DOS and NetWare). See "Restarting Failed Devices," Chapter 9, for additional information on reinitializing failed devices.

Next, if your device is listed as "Dev-Unrecognized device...", you are using a backup device that is not supported by TapeWare. In general, TapeWare supports most SCSI devices and some IDE devices. Please contact our technical support department at Support@TapeWare.Com for information on how to add support for your device.

Visit the TapeWare website at www.TapeWare.Com to see if new support for your device has been added via a service pack.

26. My autoloader under Windows NT is not recognized by TapeWare. The device is displayed, but no loader.

Make sure that you do not have any other backup programs installed. Typically, other backup programs, such as BackupExec, install a driver that is incompatible with TapeWare.

Troubleshooting the Storage Management Database

27. How do I choose my storage management database location?

The location of the storage management database can be very important. On larger systems, to minimize the time required to perform disaster recovery, we recommend the storage management database be located on a dedicated storage management server. This server performs no operations except the storage management processing. Place the devices on the servers to be backed up. If the storage manager server fails, no data is lost and recovery of the server can proceed in a non-critical manner. If any other server fails, recovery can be quickly performed since the storage server is still operational.

On smaller systems, it is acceptable to place the storage management database on the same server as the devices. Recovery does not take as much time since the databases are smaller.

For additional information, see “Managing the Storage Management Database,” Chapter 10.

28. When recovering from a storage server failure, do I get all the information from the TapeWare database back?

Most of the database is restored when you restore the storage management database. The only thing that will not be restored is the complete log of the job that was running when the database was backed up. The reason for this is that the log is not written until after the job has completed. Jobs are not completed until the database is written to the media.

Troubleshooting Error Messages

29. When attempting to add a new object or during a backup, an ‘error 51–Database corrupt’ is displayed.

Normally, if the database is corrupted, TapeWare will automatically repair it during initialization. However, it is possible that the quick check used during initialization may not detect any errors. To force TapeWare to repair the database, edit the TapeWare.Ini configuration file in the directory where you installed TapeWare. Change the line that reads **repairDatabase=No** in the [configuration] section to **repairDatabase=Yes**.

The next time TapeWare is started, the database will be automatically repaired.

30. An 'error 212-Unknown error' is returned when I restore files under Windows NT.

Windows NT stores security information in the data stream on the backup media. This information depends on the registry on the system from the files original location. If you are restoring to a different system, or are restoring a file to the same machine with a new registry, the security information contained on the media is no longer valid.

To prevent this problem, open the **Advanced Options** window from the **Options** tab of the restore job and clear the **Parent security** and **Child security** options. This causes TapeWare to restore the data in the file, but not security information such as the owner or access control lists for the file.

Disaster Recovery

The TapeWare disaster recovery system has been designed to be as automatic as possible during both initial preparation and recovery. Once installed, disaster recovery will perform its tasks without any intervention by you.

Requirements

To create the disaster recovery tapes and prepare your system for disaster, the Disaster Recovery option requires:

- For NetWare systems
NetWare 4.11 or above
- For Windows systems
Windows NT 4.x
- VGA compatible display adapter
- At least one BIOS addressable hard disk

When performing the recovery procedure, TapeWare assumes that major changes to your hardware have not occurred. The hardware you are restoring to must be nearly identical to the source system with the following exceptions:

- You may change your video adapter as long as the new video adapter is VGA compatible.
- You may increase the size of your hard disk, but the geometry of the hard disk should remain the same. For example, if your original system had a hard disk with 63 sectors per track and 255 heads, then the new hard disk should be the same. The actual number of cylinders can be larger. However, if the geometry does change, TapeWare will still use it, but the recovered operating system may not function correctly.

- Your SCSI or ATAPI tape and adapter must be the same or use the same driver as it did when the Disaster Recovery tape was created.
- You may change network cards, USB ports, and USB peripherals without restriction.

Note Ideally, you should perform the disaster recovery operation on the same computer after replacing the faulty hardware that caused the system failure.

Installing Disaster Recovery

To install the optional Disaster Recovery package, start the Installation Manager. Platform specific instructions are found in the section “Installation Instructions,” Chapter 2.

In the Installation Manager window, select Install Package. The Installation Manager will guide you through the proper installation steps. From the available packages, select the Disaster Recovery Options.

You may install or remove the Disaster Recovery option at any time without affecting the normal operation of TapeWare.

Using HP OBDR

If you are using an HP One Button Disaster Recovery (OBDR) tape drive, and your system is compatible with this feature, you do not need to make diskettes. However, we recommend that you create a set of diskettes to boot your system from in case you need to change to a non-OBDR compliant controller or tape device during recovery.

Warning Before relying on tapes being bootable with OBDR, you should check your system for compatibility as described in the OBDR Compatibility section later in this section.

Configuring Disaster Recovery

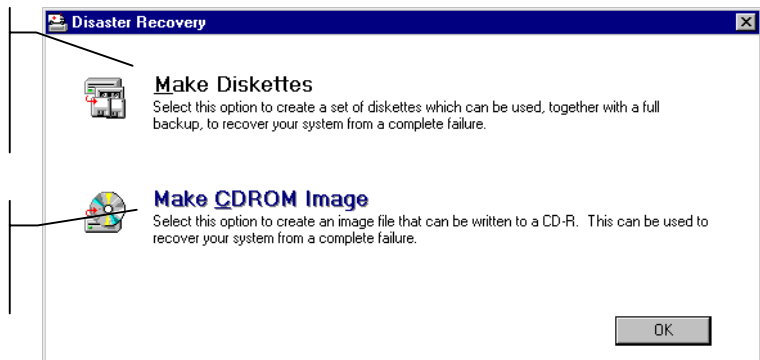
The TapeWare Disaster Recovery option configuration is completely automatic. However, you must create a set of diskettes or a CDROM to use during the actual recovery process. These diskettes, or the CDROM, are used to boot your system to initiate the recovery process after a system or disk failure.

To create these diskettes or CDROM, login to your system as the Admin user or equivalent and use one of the following procedures.

Note We recommend that you create at least two sets of diskettes, in case the first set of diskettes fails due to a read error on the floppy disk drive.

Select **Make Diskettes** to create the diskettes required to boot your system for recovery.

Select **Make CDROM Image** to create an image to write to a CD-R drive.



Windows - Diskettes

This procedure will create a set of bootable floppy diskettes that will be used to initiate a system recovery. This operation typically requires 3 blank, formatted diskettes, but may require more, depending on your system configuration.

1. Login to TapeWare as normal.
2. Select Disaster Recovery from the TapeWare Wizards tab.
3. Select Make Diskettes from the Disaster Recovery screen.
4. Follow the on-screen prompts until all diskettes are created.

Windows – CDROM

This procedure will create a file in your TapeWare main directory called BOOTIMG.DAT which can be used to create a bootable ISO-9660 CDROM. The image file created contains the entire ISO-9660 bootable image required to boot your system in the event of system failure.

Note TapeWare does not create the CDROM itself, rather it creates an image file that can be used with most CD-R mastering software and a CD-R drive to create a bootable CDROM. After TapeWare has created the image file, use your CDROM mastering software such as Nero or Easy-CD Creator to write the image directly to the CDROM. Be sure to select to image write mode, not normal file write mode.

1. Login to TapeWare as normal.

2. Select Disaster Recovery from the TapeWare Wizards tab.
3. Select Make CDROM Image from the Disaster Recovery screen.
4. After several minutes of processing, TapeWare will return to the Disaster Recovery screen.
5. Start your CD-R mastering software.
6. Select to burn an image directly to the CD-R and use the BOOTCD.IMG file created in step 3.

NetWare - Diskettes

This procedure will create a set of bootable floppy diskettes that will be used to initiate a system recovery. This operation for NetWare 4, typically requires 3 blank, formatted diskettes, and 5 diskettes for NetWare 5, but may require more, depending on your system configuration.

1. Login to TapeWare as normal.
2. Select Other Options from the Available Options menu.
3. Select Disaster Recovery from the Other Options menu.
4. Select Create Boot Disks from Disaster Recovery menu.
5. Follow the on-screen prompts until all diskettes are created.

NetWare – CDROM

This procedure will create a file in your TapeWare main directory called BOOTIMG.DAT which can be used to create a bootable ISO-9660 CDROM. The image file created contains the entire ISO-9660 bootable image required to boot your system in the event of system failure.

Note TapeWare does not create the CDROM itself, rather it creates an image file that can be used with most CD-R mastering software and a CD-R drive to create a bootable CDROM. After TapeWare has created the image file, use your CDROM mastering software such as Nero or Easy-CD Creator to write the image directly to the CDROM. Be sure to select to image write mode, not normal file write mode.

1. Login to TapeWare as normal.
2. Select Disaster Recovery from the TapeWare Wizards tab.
3. Select Make CDROM Image from the Disaster Recovery screen.

4. After several minutes of processing, TapeWare will return to the Disaster Recovery screen.
5. Start your CD-R mastering software.
6. Select to burn an image directly to the CD-R and use the BOOTCD.IMG file created in step 3.

When to Re-create the Boot Diskettes

You must re-create these diskettes under the following circumstances:

- You update your operating system by installing a service pack or other software.
- You add or remove hardware from your computer.
- You change the configuration of your disk drives, add or remove volumes or partitions.

Note You should create two sets of diskettes for backup purposes in case the first set is faulty or cannot be read by the floppy disk drive when needed for recovery.

Preparing for Disaster

After you have created your boot floppy diskettes or CDROM, you are ready in case disaster strikes. TapeWare will automatically produce Disaster Recovery Ready tapes whenever a full, overwrite backup job is run. To create a tape that is compatible with the TapeWare Disaster Recovery use the following steps:

1. Start TapeWare as normal.
2. From the TapeWare Wizard panel, select Backup Local System.
3. Follow the onscreen guides making sure that the Write type is set to overwrite. The default is append, so you must change this option to overwrite.
4. Run the job.

However, to make your disaster recovery process as easy as possible, you should note the following:

- TapeWare will rewrite most system configuration information to the tape whenever the tape is overwritten. You can select Overwrite Mode from the Backup Job Options tab.

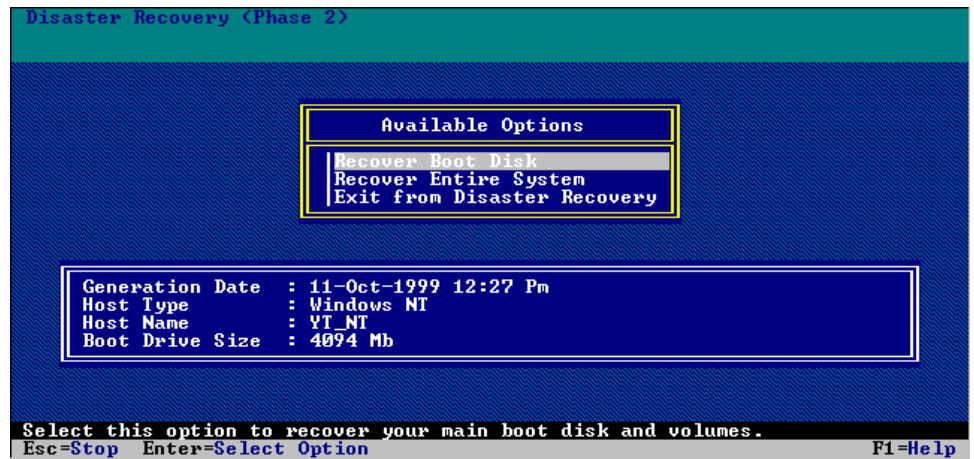
- Disaster Recovery works best with full backups, where all disks on your system fit on a single tape. If the total amount of data on your system requires more than one tape, TapeWare will still function correctly. However, you will have to change tapes in the middle of the recovery procedure. Incremental and differential jobs can also be used for recovery purposes but you must be sure to insert the recovery tapes in the correct chronological order.
- With autoloaders, you must make sure that the tape in Slot 1 contains the most recent backup. For more information, see *Using Autoloaders with Disaster Recover* later in this section.
- When recovering your system, you will be given the option to recover your whole system, or just the hard disk that your system boots from. If the volumes on your boot hard disk are split among multiple physical hard disks, whether using mirroring, volume extension or striping, TapeWare may not be able to correctly recover the volumes on the boot disk. You will have to restore them manually after the initial operating system is recovered. For maximum flexibility, keep the volumes and partitions on the boot disk separate from any other volumes on other disks.
- The Disaster Recovery Option must be installed on each system for which recovery is to be available. The recovery information that is written to the tape is from the local system only. No remote disaster recovery is supported. For example, if the tape device is connected to Machine1 and you make a remote backup of Machine2, the tape will contain system configuration information from Machine1. The tape created can be used to boot Machine1, but if it is used during the actual recover process, no files will be restored since all files on the tape are from Machine2.

Recovering from a Disaster

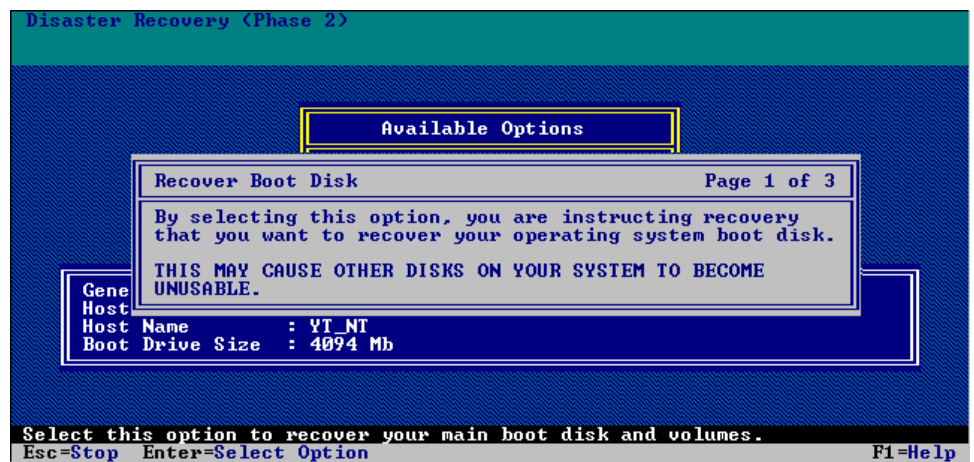
If a disaster strikes and you are unable to boot your system using your normal boot procedure, you may use the following procedure to recover your system.

Note The disaster recovery procedure can cause data loss in cases where not all disks need to be recovered. You should attempt all less-intrusive recovery procedures prior to attempting to use the Disaster Recovery procedure. For example, attempt to boot DOS from a floppy disk and check your operating system files and partitions. If you have had to replace the hard disk due to a hardware failure, you must use the following procedure to recover the disk.

1. If you have a One Button Disaster Recovery (OBDR) compatible tape drive:
 - A. Insert the tape you want to boot from into the tape drive.
 - B. Press and hold down the Eject button on the front panel of the tape drive while turning the power on to the tape drive. After a short delay (2 seconds), the lights on the drive will start blinking back and forth. Once the lights start blinking, release the Eject button.
 - C. Turn the power on to your computer.
 - D. After your normal BIOS startup and system diagnostics are executed, your computer should start booting from the tape in the tape drive.
 - E. If your system does not boot properly from tape, due to a change in hardware or an incompatibility, proceed to Step 2 to boot from the Disaster Recovery floppy diskettes or CDROM.
2. If you do not have an OBDR compliant drive, you must use the boot floppies or CDROM that you created in the preparation steps.
 - A. Insert the first Disaster Recovery diskette into your floppy drive.
 - B. Turn on the power to your computer.
3. After your system has booted, either from tape or diskette, select Recover Boot Disk or Recover Entire System from the menu. If you suspect that only your boot disk is corrupted or you have had to replace the boot disk, you may attempt to recover the boot disk only. This will leave all data on other disks intact. However, this may cause data to be overwritten on other disks as well.



4. After selecting one of the two recovery options, you must confirm several times that this is really what you want to do. If you are absolutely sure, you may proceed to step 5.



5. If you booted your system from floppy diskettes, you are requested to insert the rest of the diskettes for files to be copied onto to the boot hard disk.
6. There is no further interaction required by you until after the first tape is restored to your system. This entire process of restoring the first tape can be as little as 15 minutes or as long as 2-3 hours, depending on how much data is contained on this tape, the speed and capabilities of the tape drive, and whether you are recovering the whole system or just the boot disk.

7. After the first tape is restored, you are asked if there are any more tapes that you need to restore. If you have any more tapes, select “Yes” or press F10. You may require multiple tapes if your system required more than one tape to complete a full backup, or if you need to restore any additional incremental or differential tapes. However, you should use full backups only for recovery, and restore incremental and differential tapes after your system is running again.

Note You should recover from full backups only. After the recovery is completed and your system has been restarted, use the standard TapeWare restore options to restore any additional tapes, incremental or differential, to your system. The standard TapeWare restore procedure optimizes the restoration process and can be much faster restoring these incremental and differential tapes than the disaster recovery process.

OBDR Compatibility

To test your hardware compatibility with the One Button Disaster Recovery specification, perform the following steps. This will not cause you to lose any data on your system and is completely safe.

1. Make a full backup of your system as instructed in the Preparing for Disaster section.
2. Shut down your system as normal.
3. Follow the instructions in Recovering from a Disaster. However, do not select to recover.
4. If your system boots and displays the Disaster Recovery Phase 2 screen, your hardware and BIOS are compatible with OBDR. If this screen is not displayed and your normal operating system is booted, or your system hangs, you must use floppy-diskette based disaster recovery. See the Configuring Disaster Recovery earlier in this section.
5. Reset your computer by pressing the reset button on the front. This will end the disaster recovery procedure and end the compatibility test.

Using Disaster Recovery with Autoloaders

When using an autoloader with the TapeWare Disaster Recovery option, remember the following:

First, you must know which tape in the autoloader contains your most recent full backup. This is important because TapeWare will boot only from the tape in slot 1 of your autoloader. So, prior to starting the recovery process, you should ensure that this most recent full backup is contained in slot 1.

Next, TapeWare will restore all tapes that are contained in the autoloader during the final recovery process. So, you need to make sure that the tapes in your autoloader at the time of recovery are tapes that you actually need to restore. For example, if you are using one of the standard Simple 6 slot rotation schedules, TapeWare will have a monthly full backup on one of the tapes, and differential backups on the other tapes. All you should restore from the disaster recovery process is the full backup tape. Then, after your system is recovered, use the standard TapeWare restore procedures to recover any more recent data from the differential backup tapes.

When using an autoloader, we recommend removing from the autoloader slots, all tapes except the full backup contained in slot 1. If the full backup spans more than a single tape, put the additional full backup tapes into additional slots. TapeWare will then restore these tapes along with the tape in slot 1.

Note It is critical that slot 1 contain the most recent full backup that was written in overwrite mode. When using overwrite mode in the backup options, TapeWare writes key system configuration information to the beginning of the backup tape that is used to reconfigure your system.

Troubleshooting – Preparing for Disaster

When creating a set of disaster recovery diskettes or when creating the boot tracks written to tape at the beginning of a backup, the following problems may occur:

1. File not Found error:

Examine the TWTrace.Txt file located in your TapeWare installation directory. This will list the file that was not found. All files must reside in “standard” directory locations provided by the operating system. For example, under Windows NT, all drivers must reside in the \WINNT\SYSTEM32\DRIVERS subdirectory. For NetWare, drivers must reside in either SYS:SYSTEM or C:\NWSERVER. If the file is not a driver file for your hardware, please contact technical support for further assistance.

2. Registry read or write error:

TapeWare requires full access to the registry to create disaster

recovery information. Make sure you are starting TapeWare with an account with full administration privileges.

3. Get/Put Server Information failure:

An error has occurred while attempting to retrieve your disk configuration information. Consult the TWTrace.Txt file for more specific details on this error. Disks that are not currently powered on or a lack of security rights usually cause this error.

Troubleshooting – Recovering from a Disaster

When attempting to recover your system, the following are common errors that occur:

1. Get/Put Server Information failure:

This error is caused by a failure to reconfigure your disk drives and volumes on the target system. Make sure that all disks are powered on and ready, and that any new disks are of the same size or larger. Also, make sure that the geometry's of any new hard disks are the same as the old disks. Check to make sure that you have enabled logical block addressing in your BIOS configuration, and that any SCSI controllers are configured in the same way they were when the recovery tape was created. If you have changed SCSI adapters, TapeWare may not be able to access any peripherals on the new adapter unless the new adapter is using the same driver as the old adapter.

2. Missing DISPLAY_DRIVER.DLL on NT start:

When starting TapeWare immediately after the initial recovery screen, TapeWare switches into Windows NT mode. If you have changed video adapters and your new adapter is not fully VGA compatible, this error is displayed.

3. No Tape Drives or Controllers present:

If you have changed SCSI controllers, this message may appear when the new controller is not compatible with the older controller driver, and your tape device was connected to the older controller. TapeWare was unable to locate the tape device on the new controller.

4. Boot track not found:

When attempting to boot from tape, if the tape does not contain a valid disaster recovery boot track, this message will be displayed. Try another tape, or another tape drive, if available.

Index

A

- Access permission, 49, 201
- Access range
 - backup jobs, 65
 - restore jobs, 78
- Address tab, 204
- Admin folder, 39
- Advanced options
 - restore jobs, 126
- Advanced Options
 - backup and restore jobs, 120
- Arranging objects, 22
- Audit tab, 205
- Auto format mode, 237
 - backup jobs, 112
- Auto verify mode, 236
 - backup jobs, 116, 236
- Autoloaders
 - cleaning, 154
 - drivers, 153
 - loader test tab, 224
 - setting up for automatic cleaning, 182
 - status tab, 253

B

- Backup devices
 - and streaming, 167
 - installing, 13
 - locating on machines, 168
 - troubleshooting, 298
- Backup job
 - options tab, 109
- Backup jobs

- advanced options, 120
- and manual schedule, 104
- and selection filters, 63
- backup mode, 110
 - creating by copying, 36
 - creating from job tab, 34
 - creating with wizard, 34
 - options automatically updated, 119
 - purposes, 32
 - Renaming, deleting, and moving, 37
 - rotating media, 92
 - scheduling concepts, 91
 - scheduling with built-in schedule, 99
 - selecting a schedule, 95
 - selecting files, 62
 - troubleshooting, 293
- Backup mode, 93, 234
 - and forcing scheduled jobs, 173
 - and Microsoft Exchange server, 270
 - and Microsoft SQL server, 277, 278
 - and Microsoft Exchange server, 272
 - backup jobs, 110
- Backup range
 - restore jobs, 75
- Backup Range
 - backup jobs, 64
- Backup stream, 255
- BTrieve
 - configuring database sets, 261

C

- Century date algorithm, 63
- Change mode, 119, 235
- Child security, 122

- Children
 - backup jobs, 67
- Clean device, 154
- Cleaning Autoloaders, 182
- Command tab, 206
- Communication test tab, 208
- Connections tab, 210
- Containers
 - concepts, 4
 - display, 20
- Copying
 - directory structures, 178
 - to create new jobs, 36
- Create permission, 49, 200
- Create range
 - restore jobs, 76
- Create Range
 - backup jobs, 65
- Creating job folders, 40
- Creating jobs
 - by copying, 36
 - from job tab, 34
 - with wizard, 34
- Creating new media, 149
- Creating new users, 186
- Custom schedule, 100
 - modifying, 102
- Customer support, xx

D

- Data streams
 - storage tab, 255
- Data streams
 - optimal number, 168
- Database
 - restoring, 158
- Database tab
 - and folders, 38
 - and jobs, 33
- Date
 - 20th vs 21st century, 63
- Delay between passes
 - restore jobs, 125
 - verify jobs, 128

- Delay Between passes
 - backup jobs, 115
- Delete permission, 50, 199
- Delete range
 - restore jobs, 76
 - using to select deleted files, 175
- Deleted files
 - selecting for restoring, 175
- Deleting folders, 41
- Deleting jobs, 37
- Deleting media, 152
- Device
 - options tab, 238
 - options tab, backup jobs, 117
 - options tab, restore jobs, 125
 - verify jobs, 129
- Device tab, 152
 - compared to media tab, 147
- Devices
 - installing, 13
 - restarting failed, 153
 - troubleshooting, 298
- Diagnostics tab, 211
- Documentation
 - additional, xix
- Drivers tab, 212

E

- Effective permissions, 192
 - abbreviations, 217
 - algorithm, 193
 - and equivalencies, 192
 - and groups, 192
 - calculating, 45, 192
 - checking, 171, 196
 - concepts, 44
 - examples, 45, 194
 - from multiple sources, 193
 - how moving effects, 47
 - overview, 43
- Eject magazine, 154
- Eject media, 154
- Email
 - configuring protocol, 265

- installing, 265, 301
- Email tab, 213, 267
- Equivalencies, 188
- Equivalencies tab, 214
 - and assigning permissions, 189
- Erase quick, 155
- Erase secure, 155
- Errors
 - stream sync error, 297
 - troubleshooting, 299
- Everyone folder, 39
 - and permissions, 40
- Everyone group, 44, 188
- Exclude attributes
 - backup jobs, 67
 - restore jobs, 79
- Expiration date, user accounts, 187

F

- Filters. *See* Selection filter
- Find button, 23
- Folders
 - admin folder, 39
 - creating, 40
 - creating media, 148
 - deleting media, 149
 - different types, 38
 - everyone folder, 39
 - home folder, 39
 - job, 37
 - moving, renaming, and deleting, 41
 - new user/group, 186
- Format media window, 155

G

- General tab, 216
- Grace logons, 17, 187
- Groups
 - creating new, 191
 - using for complex security, 171
- Groups tab, 188, 220

H

- Help
 - on-line, xx
- Home folder, 39

I

- Identify media, 157
- Import media, 157
- Incremental jobs
 - and data recover period, 104
- Installation, 9
 - and number of database zones, 10
 - and storage management server, 163
 - as a Windows service, 13
 - backup devices, 13
 - before you install, xix
 - configuring BTrieve database sets, 261
 - email, 265, 301
 - installing database server first, 10
 - instructions, 11
 - multiple platforms, 10
 - number of servers, 10
 - product configuration, 10
 - troubleshooting, 291
- Installation manager window, 12
- Installing options (packages), 12
- Instance range
 - and selecting files not backed up previously, 174
 - backup jobs, 65
 - restore jobs, 78
- Instance window, 222
- Instances
 - latest instance, 70
 - selecting folders by date, 70
 - selecting for restore jobs, 69, 73
 - selecting from specific job, 175
 - selecting from specific media, 176
 - selecting specific instances, 74
- Instances and filters, 70
- Instances window, 70

J

Job folders, 38
Job status window, 139
Jobs. *See* Backup jobs, Restore Jobs, Verify jobs

K

Keyboard Shortcuts, 22

L

Loader test tab, 224
Log option
 backup jobs, 115
 restore jobs, 125
Log options
 verify jobs, 128
Logging on
 default password, 17
 grace logons, 17
 instructions, 17
 troubleshooting, 292
Logging out, 17
Logon control tab, 186, 225
Logon window, 14
logons
 grace, 187
Logs
 emailing, 265, 301
 viewing and printing, 141
Logs tab, 228

M

Macintosh Finder, 123
Manual rotation schedule, 104
Manual schedule
 and verify jobs, 106
 restore jobs, 105
Marking files, 58
Media
 backup job options, 118
 creating new, 149
 deleting, 152
 formatting, 149, 155

 identifying, 157
 importing, 157
 importing and security, 184
 number required, rotation jobs, 98
 options tab, 239
 password, options tab, 118
 passwords, 151
 passwords and security, 184
 selecting instances from, 176

Media control tab, 229

Media filter

 backup jobs, 67
 restore jobs, 80

Media folders, 38

 creating, 148
 deleting, 149

Media password, 239

Media tab, 148

 and folders, 38
 compared to device tab, 147

Members tab, 191, 231

Menus, 25

Microsoft Exchange, 269

 and backup modes, 272
 configuration window, 270
 restoring databases, 273

Microsoft SQL Server, 275

 and backup modes, 277, 278
 configuration window, 276
 restoring databases, 279

Modify permission, 50, 199

Modify range

 restore jobs, 76

Modify Range

 backup jobs, 64

Moving folders, 41

Moving jobs, 37

Moving objects

 and permissions, 47

N

Native data streams format, 121

 and transferring between platforms, 181

New media location, 237

- backup jobs, 113
- New media name, 238
 - and forcing scheduled jobs to run, 173
 - backup jobs, 112
- NFS information, 123
- Number of concurrent connections, 187
- Number of passes
 - backup jobs, 114
 - restore jobs, 124
 - verify jobs, 127

O

- Object owner, 123
- Objects
 - concepts, 4
- Options tab, 233
 - and running failed jobs, 172
 - backup jobs, 109
 - backup mode concepts, 93
 - log option, 228
 - options automatically updated, backup jobs, 119
 - restore jobs, 123
 - verify jobs, 127
- OS/2 extended attributes, 122
- Owner, job, 136
 - and email support, 268

P

- Parent security, 122
- Parents
 - backup jobs, 67
 - restore jobs, 79
- Passwords, 187, 226
 - and importing media, 157
 - changing, 16
 - changing when forgotten, 16
 - media, 239
- Permissions
 - access, 49, 201
 - and everyone group, 44
 - and forcing jobs to run, 136
 - checking effective permissions, 171
 - create, 49, 200

- delete, 50, 199
- direct. *See* Effective permissions
- effective, 192
- examples, 50
- granting to other users, 53
- inherited. *See* Effective permissions
- modify, 50, 199
- overview, 43. *See also* Effective permissions
- planning considerations, 183
- read, 50, 198
- restricting, 54
- supervisor, 48, 201
- tips and strategies, 170
- write, 50, 198

- Permissions tab, 241

- groups, 191
 - users, 189

- Ping test tab, 243

- Preferences window, 244
 - smart expandable indicators, 21

- Properties
 - concepts, 5
 - property sheets, 23

- Property sheets
 - opening, 24

Q

- Query filters window, 161
- Query window, 257
- Queue tab, 132

R

- Read permission, 50, 198
- Renaming files
 - during restore job, 81
- Renaming folders, 41
- Renaming jobs, 37
- Reports, 26
 - printing under NetWare, 295
- Required attributes
 - backup jobs, 66
 - restore jobs, 79
- Restore database command, 158

- Restore jobs
 - advanced options, 120
 - concepts, 32
 - creating by copying, 36
 - creating from job tab, 34
 - creating with wizard, 34
 - options tab, 123
 - renaming, deleting, and moving, 37
 - scheduling, 105
 - selection concepts, 68
 - troubleshooting, 295
- Restore Jobs
 - selecting instances, 73
- Restoring
 - files to new folder, 81
 - for a specific date, 176
 - Microsoft Exchange database, 273
 - Microsoft SQL databases, 279
 - to new or different folder, 73, 80, 179
 - with latest date, 176
 - with new names, 81, 180
- Retention, 154
- Rewind, 154
- Rotation media, 92
- Rotation set, 92
- Running jobs
 - and logging out, 18
 - automatically running scheduled jobs, 133
 - concepts, 7
 - forcing and job options, 135
 - forcing scheduled jobs, 134
 - running failed jobs again, 172
 - scheduling to run once only, 173
 - security and scheduled jobs, 134
 - strategies for faster, 167
 - unscheduled jobs, 138

S

- Schedule tab, 246
- Scheduled jobs
 - running only once, 173
 - tips and strategies, 172
- schedules

- and full data recovery period, 103
- Schedules
 - and full data recovery period, 97
 - and restore jobs, 105
 - and verify jobs, 106
 - backup job types, 93
 - built-in schedules compared, 98
 - built-in, backup jobs, 96
 - concepts, for backup jobs, 95
 - customizing, 100
 - modifying built-in, 101
 - modifying custom schedule, 102
 - not scheduled, backup jobs, 96
 - rotation schedules compared, 97
 - which to select, 97
- Scheduling jobs
 - concepts, 91
- Security
 - and logging out, 18
- Security tab
 - adding new users and groups, 185
- Select Instance button
 - and selecting instances, 74
- selecting files
 - instances, verify jobs, 84
- Selecting files
 - automatically selecting new files, 61
 - changed files only, 61
 - concepts, 58
 - deleted, for restoring, 175
 - files vs. folders, backup jobs, 59
 - files vs. folders, restore jobs, 72
 - files vs. folders, verify jobs, 84
 - filter criteria, verify jobs, 86
 - for backup jobs, 58, 62
 - for restoring, 68
 - for verify jobs, 83
 - grayed or checked boxes, 58
 - instances of files, restore jobs, 69
 - instances on specific media, 176
 - instances with specific date, 175
 - not previously backed up, 174
 - selecting instances with folders, 72
 - sorting with filters, 60
 - specific instances, restore jobs, 74

- specific instances, verify jobs, 86
- tips and strategies, 174
- with filters, verify jobs, 85
- Selection filters
 - and grayed folders, 61
 - criteria, backup jobs, 63
 - criteria, restore jobs, 74
 - including vs. excluding, 60
 - selecting files with, 60
- Selection filters window, 249
- Selection tab, 248
- Service, Windows, 13
- Set count, 92
 - and custom schedule, 102
- Set name, 92
- Shortcut Menus, 23
- Size range
 - backup jobs, 65
 - restore jobs, 78
- Space restrictions, 122
- Starting TapeWare, 14
- Status bar, 21
- Status field
 - instance window, 83
- Status tab, 253
- Storage management database
 - accounting for size, 166, 167
 - object hierarchy illustration, 197
 - repairing corrupt, 299
 - restoring, 158
 - troubleshooting, 298
 - where to locate, 163
- Storage Management Database
 - concepts, 3
- Storage management server, 163
 - and installation, 10
 - selecting, 14
- Storage Management zone
 - selecting, 14
- Storage management zones
 - and backup devices, 6, 164
 - and machines, 6
 - and security issues, 183
 - remote administration of jobs, 165
 - security, 6

- storage management server, 164
- Storage tab, 255
 - using to increase job speed, 169
- Supervisor permission, 48, 201

T

- Toolbar, 19
- Transferring data between operating systems, 181
- Troubleshooting
 - backup devices, 298
 - backup jobs, 293
 - Errors, 299
 - installation, 291
 - restore jobs, 295
 - storage management database, 298
 - verifying, 297

U

- Updating license, 12
- Upgrading
 - updating license, 12
- User/group folders, 38
- Users
 - adding new, 186
 - controlling logon, 187
 - logon control tab, 186

V

- Verify jobs
 - concepts, 33
 - creating by copying, 36
 - creating from job tab, 34
 - creating with wizard, 34
 - options tab, 127
 - Renaming, deleting, and moving, 37
 - scheduling, 106
 - selecting files, 85
 - troubleshooting, 297
- Verify mode
 - verify jobs, 128
- Volume restrictions, 122

W

Wildcard format, 66

Wildcard type
 restore jobs, 79

Wizard
 overview, 28

Write mode, 235

and forcing scheduled jobs to run, 173
 backup jobs, 111

Write permission, 50, 198

Z

Zones. *See* Storage management zones