



iSeries

Optical Support

Version 5

SC41-5310-04





@server

iSeries

Optical Support

Version 5

SC41-5310-04

Note

Before using this information and the product it supports, be sure to read the information in "Notices" on page 173.

Fifth Edition (April 2004)

This edition applies to version 5, release 3, of IBM Operating System/400 (product number SC41-5310-04) and to all subsequent releases and modifications until otherwise indicated in new editions. This edition applies only to reduced instruction set computer (RISC) systems.

This edition replaces SC41-5310-03.

© Copyright International Business Machines Corporation 1997, 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	vii
--------------------------	------------

About Optical Support (SC41-5310). . . . **ix**

Who Should Read This Book.	ix
Conventions and terminology used in this book	ix
Prerequisite and related information	x
iSeries Navigator	x
How to send your comments.	x

Summary of Changes **xiii**

Chapter 1. Optical Support—Introduction **1**

Optical Media Types.	1
Optical System Configurations	2
Directly-Attached Libraries	2
LAN-Attached Libraries	3
Supported Hardware for Optical Storage	4
Optical Storage Organization	7
Optical Volumes	7
Optical Directories	8
Optical Files	8
Using and Managing Optical Storage	8

Chapter 2. CD-ROM and DVD on iSeries Servers **9**

Supported Media Formats	9
Configuring Your CD-ROM or DVD Drive	9
Working with Volume Identifiers	10
Loading and Unloading CD-ROM and DVD Media	10
Duplicate Volume Identifiers.	11
Allocating the Device Description	11
Non OS/400 access of optical devices	12
Virtual Optical Devices	12
Mastering Your Own CD-ROM.	13

Chapter 3. Directly-Attached Optical Media Libraries **15**

Connectivity of non-IBM library devices.	15
Configuring Optical Media Libraries	15
Labeling Optical Cartridges (WORM or Erasable).	15
Labeling a New Optical Cartridge with Uninitialized Volumes	16
Labeling an Optical Cartridge with an Initialized Volume	16
Setting Write Protection	17
Allocate/Deallocate Optical Drive	17
Getting Started with Optical Cartridges and Volumes	19
Adding Optical Cartridges to an Optical Media Library—Example	19
Initializing Optical Volumes—Example	20
Removing an Optical Disk Cartridge—Example	22
Changing Optical Environment Parameters.	23

Chapter 4. LAN-Attached Optical Media Libraries **25**

Configuring Optical LAN Support.	27
OS/400 Communications Control Setup—Example.	29
LAN Controller Setup	30
Configuring Optical LAN Support over a TCP/IP connection.	30
Volume Security	31
Description of LAN-Specific Commands.	32
Add Optical Server.	32
Remove Optical Server	32
Display Optical Server.	32
Differences between Directly-Attached and LAN-Attached Devices	33
Getting Started with Installing an Optical LAN Device	34
Adding an Optical LAN Device	34
Working with Volumes in a LAN Device	35
Media Interchange Between LAN and Direct Attach Libraries	36
Removing an Optical LAN Device.	36

Chapter 5. Working with Optical Volumes **37**

Command and Device Dependency	37
Working with Optical Volumes Option	38
Displaying Optical Volumes	38
Initializing Optical Volumes	40
Renaming Optical Volumes	41
Adding Optical Disk Cartridges	41
Copying Optical Volume Data	42
Changing Optical Volume Attributes	44
Displaying and Printing Optical Volume Attributes	45
Duplicating Optical Volumes	47
Working with Directories and Files	48
Removing Optical Volumes	49
Check Optical Volume.	50

Chapter 6. Optical Security and Auditing **53**

Authorization List Security	53
Authorities Required for Optical Functions	54
Specifying an Authorization List on the Add Optical Cartridge (ADDOPTCTG) Command.	55
Authorization List Assignment and the Add Optical Server (ADDOPTSRV) Command	55
Changing the Authorization List to Secure an Optical Volume	55
Retaining the Authorization List when Removing Optical Volumes.	56
Mapping an Authorization List to an Optical Volume.	56

Directory and File Level Security	56
Optical Auditing.	56
Object Authority.	57

Chapter 7. Using Application Programming Interfaces for Optical Support 59

Categories of APIs for Optical Support	59
Using the HFS API Interface to Optical Support	59
Optical File System Implementation of HFS APIs	60
Control File System Functions to Optical Support	67
Standard Attribute Definition and Restrictions—Directly-Attached.	71
Special Attributes Unique to Optical File System	72
Copying File Attributes Using HFS	74
Using the Integrated File System Support Interface to Optical Support	75
Optical File System Implementation of Integrated File System Support	75
Thread Safety.	82
Tips and Techniques for the Application Programmer	82
Media Capacity and Volume Threshold	82
Expanding Buffer I/O through HFS—Tailoring Read Requests for Performance.	84
Force Buffered Data or fsync() APIs	85
Held Optical Files	85
Volume, Directory, and File Considerations.	85
Path Names	87

Chapter 8. Optical Volume Backup 89

Defining Your Backup Strategy	89
Your Backup Options	89
DUPOPT	90
Enhancements	91
CPYOPT	91
Key parameters	92
Using CPYOPT to Backup a volume	94
Volume type *BACKUP	96
CPYOPT Performance	100
SAV/RST.	102
SAV Support	103
RST Support.	104

Chapter 9. Optical Media Library Performance. 107

General Performance Considerations for Optical Support	107
Volume Mounting and Dismounting.	107
Drive Contention	107
Number of Directories and Files	107
File Size	107
Performance Considerations—Directly-Attached Optical Libraries	107
Effects of File Attributes on Performance	107
HFS API Expanding Buffer I/O	108
CPYOPT and DUPOPT Performance Considerations	108
Volume Mounting and Dismounting Schedule Timers.	108

Add Optical Cartridge Performance	109
Performance Considerations—LAN-Attached Optical Libraries	109
Controlling the Block Size for an Optical LAN Server	109
Conversation Allocation for Optical LAN Server	110
Minimizing Volume Mounts for Optical LAN Server	111

Chapter 10. Optical Media Formats 113

ISO 9660	113
Overview.	113
Volume, Directory, and File Names	114
Programming Interfaces	114
Command Interfaces	114
Directory and File Security	114
High Performance Optical File System	114
Overview.	114
Volume, Directory, and File Names	115
Space Reclaim	115
Programming Interfaces	115
Command Interfaces	116
Directory and File Security	116
Media Interchange.	116
Directory Structure and Performance	117
Universal Disk Format	117
Overview.	117
Volume, Directory, and File Names	117
Programming Interfaces	118
Command Interfaces	119
Directory and File Security	119
Media Interchange.	119
Directory Structure and Performance	120
Command and Media Format Dependencies	120
Differences Between HPOFS and UDF	121

Appendix A. Reclaiming the Optical Index Database 123

Optical Index Information	123
Optical Index Database Files	123
Internal Library Index	124
Reclaiming the Optical Index for a Stand-Alone Optical Device	124
Reclaim Types	124
*SYNC—Synchronize Volume Index with Internal Library Index	124
*UPDATE—Update Volume Index from Optical Volumes	125
*RESET—Reset Internal Library Index and Reclaim Volume Index	125
Time Required to Complete Reclaim Optical Index	126
Choosing the Reclaim Type to Use	126

Appendix B. Optical Save and Restore 129

Introduction.	129
Save and Restore Command Summary	129
Operational Hints	129
Optical Media Formats	129

Organizing Save and Restore Data on Optical Devices	130
Performing a Save Operation to Optical Storage	130
Displaying Save and Restore File Information on Optical Storage	130
OPTFILE Parameter	130
Media Eject Option	131
Volume List Processing	131
Volume Lists with HPOFS Format Media	132
Volume Lists with UDF Format Media	132
Performing a Save Operation to DVD-RAM, UDF, and HPOFS Media	133
Specifying the path for your files	133
Using the CLEAR parameter	133
Using multiple volumes	134
Command Restrictions	134
Operational Differences by Device Types	135
Optical Library Dataservers	135
CD-ROM and DVD-ROM Standalone Optical Drive Devices	135
DVD-RAM Standalone Optical Drive Devices	135

Appendix C. Recovering Held Optical Files 137

How Optical Files Are Manipulated	137
Held Optical Files	137
Recovery for Held Optical Files	138

Appendix D. Output File Support Structures. 143

Output File Structure for Volume Attributes	143
Output File Structure for Directory Attributes	145
Output File Structure for File Attributes	145

Appendix E. Handling Problems That Are Related to Optical Support. 147

Common Optical Support Problems and Questions	147
Handling Installation Problems of LAN Optical Media Libraries	150
PC-Controller Software	
Installation/Configuration Problems.	152
Collecting Information	153
Other System Commands	153

Appendix F. Programming Examples 155

HFS API Program Examples Using ILE RPG for OS/400	155
Getting a Path and Calling Subroutines.	155
Defining Data Structures for Opening Files	156
Opening an Optical File	156
Writing a File to an Optical Disk	156
Closing an Optical File	157
Integrated File System Program Examples Using ILE C for OS/400	157
Integrated File System Function Examples.	157
Optical Tools	160
Copy Stream File	160
Copy Database File to Optical File	162
Copy Spooled File to Optical	165
Copy Optical to Database	167

Bibliography. 171

Notices 173

Programming Interface Information	175
Trademarks	175
Terms and conditions for downloading and printing publications	176
Code disclaimer information	176

Index 177

Figures

1. Example of Using an iSeries server Directly-Attached Library through a Local Area Network	3	16. Work with Optical Volumes Display, Extended Information	40
2. Example of Using iSeries servers with a LAN-Attached Optical Library	4	17. Initialize Optical (INZOPT) Display	41
3. Work with Optical Volumes Display	16	18. Add Optical Cartridge (ADDOPTCTG) Display	42
4. Work with Media Library Status	18	19. Copy Optical (CPYOPT) Display	43
5. Work with Media Library Status	18	20. Change Optical Volume (CHGOPTVOL) Display	45
6. Work with Optical Volumes Display	20	21. Display Optical Volume Attributes Display, First Display	46
7. Work with Optical Volumes Display	21	22. Display Optical Volume Attributes Display, Second Display	46
8. Work with Optical Volumes Display	22	23. DUPOPT Panel	47
9. Typical Optical LAN Configuration	26	24. WRKLNK Panel	48
10. Optical LAN Software Interfaces	28	25. WRKOPTDIR Panel	49
11. Optical LAN Configuration	29	26. Remove Optical Cartridge (RMVOPTCTG) Display	49
12. Add Optical Server	35	27. Check Optical Volume (CHKOPTVOL) Display	50
13. Work with Optical Volumes Display	35		
14. Work with Optical Volumes Display	38		
15. Work with Optical Volumes Display, View 2	39		

About Optical Support (SC41-5310)

This book serves as a user's guide and reference for IBM® optical support on the OS/400 operating system. This book explains the functions that are unique to optical support, and the information contained in this book can help you use and understand:

- CD-ROM devices
- DVD devices
- Directly attached optical media library devices
- LAN-attached optical media library devices

Unless otherwise stated, this book refers to the operational and functional characteristics of directly attached optical devices. Many of these characteristics are the same for LAN-attached optical media library devices, but are not guaranteed unless stated.

Who Should Read This Book

The information in this book is intended for the following audiences:

System operators

iSeries server operators can use this book as their primary reference for CD-ROM, DVD, and optical media library operation.

Service representatives

Service representatives can use this book to perform activities as directed by the *IBM 3995 AS/400 Optical Library Dataserver: Maintenance Information* book.

Application programmers

Programmers can use this book to assist in the development of applications that use the hierarchical file system and integrated file system application programming interfaces to optical support.

End users

End users can use this book as their primary reference for information on operating and using CD-ROM, DVD, and optical media libraries.

All audiences should be familiar with the iSeries server.

Conventions and terminology used in this book

This book includes examples and displays that specify dates. The date format used is MM/DD/YY where MM represents the month, DD represents the day, and YY represents the year. A different date format may be used by your system.

Prerequisite and related information

Use the iSeries Information Center as your starting point for looking up iSeries technical information.

You can access the Information Center two ways:

- From the following Web site:
<http://www.ibm.com/eserver/series/infocenter>
- From CD-ROMs that ship with your IBM Operating System/400 order:
iSeries Information Center, SK3T-4091-04. This package also includes the PDF versions of iSeries manuals, *iSeries Information Center: Supplemental Manuals*, SK3T-4092-01, which replaces the Softcopy Library CD-ROM.

The iSeries Information Center contains advisors and important topics such as Java™, TCP/IP, Web serving, secured networks, logical partitions, clustering, CL commands, and system application programming interfaces (APIs). It also includes links to related IBM Redbooks™ and Internet links to other IBM Web sites such as the Technical Studio and the IBM home page.

With every new hardware order, you receive the *iSeries Setup and Operations CD-ROM*, SK3T-4098-02. This CD-ROM contains IBM @server IBM e(logo)server iSeries Access for Windows and the EZ-Setup wizard. iSeries Access Family offers a powerful set of client and server capabilities for connecting PCs to iSeries™ servers. The EZ-Setup wizard automates many of the iSeries setup tasks.

For other related information, see the “Bibliography” on page 171.

iSeries Navigator

IBM iSeries Navigator is a powerful graphical interface for managing your iSeries servers. iSeries Navigator functionality includes system navigation, configuration, planning capabilities, and online help to guide you through your tasks. iSeries Navigator makes operation and administration of the server easier and more productive and is the only user interface to the new, advanced features of the OS/400 operating system. It also includes Management Central for managing multiple servers from a central server.

You can find more information on iSeries Navigator in the iSeries Information Center and at the following Web site:

<http://www.ibm.com/eserver/series/navigator/>

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this book or any other iSeries documentation, fill out the readers’ comment form at the back of this book.

- If you prefer to send comments by mail, use the readers’ comment form with the address that is printed on the back. If you are mailing a readers’ comment form from a country other than the United States, you can give the form to the local IBM branch office or IBM representative for postage-paid mailing.
- If you prefer to send comments by FAX, use either of the following numbers:
 - United States, Canada, and Puerto Rico: 1-800-937-3430
 - Other countries: 1-507-253-5192
- If you prefer to send comments electronically, use one of these e-mail addresses:

- Comments on books:
RCHCLERK@us.ibm.com
- Comments on the iSeries Information Center:
RCHINFOC@us.ibm.com

Be sure to include the following:

- The name of the book or iSeries Information Center topic.
- The publication number of a book.
- The page number or topic of a book to which your comment applies.

Summary of Changes

The following items are revised with this release of this guide:

- Added support for 399F device type.
- Enhancements to Add Optical Cartridge (ADDOPTCTG) and Remove Optical Cartridge (RMVOPTCTG) to support a bulk media loader for certain 399F models.
- Added command Check Optical Volume (CHKOPTVOL) to verify the integrity of the directory structure on a optical volume and output a count of the number of directories and files on a volume.
- Added support to create writable virtual optical volumes.
- Increase the number of virtual devices that can exist in a partition from 1 to 35.
- Add support for enhanced media formats which includes UDF-bridge and Rock Ridge file and directory names.
- Enhance UDF (Universal Disk Format) support to enable full media interchange between i/p Series for media formatted with UDF version 2.01.

A vertical line (|) to the left of the text indicates a change or addition.

Chapter 1. Optical Support—Introduction

Optical storage on the iSeries server provides an economical and efficient way to store and retrieve large amounts of information at a high performance level. Optical storage devices offer significant advantages over other high-capacity storage devices, such as tape and microfilm, with faster access times and a hierarchical-type file organization. OS/400 optical storage uses files that are stored in directories and files that are stored in subdirectories similar to UNIX** or PC-based file systems. Compact disk read-only memory (CD-ROM) drives are standard equipment on iSeries Advanced Series servers that use PowerPC® technology. They provide an economical means for wide-scale distribution of programs and data. DVD-RAM drives are available as an orderable feature on many iSeries servers that are running at least Version 4 Release 5 level of OS/400. You can use the DVD-RAM for many day-to-day save and restore operations. In addition, you can use file system interfaces such as the integrated file system and hierarchical file system to store and retrieve data files on DVD-RAM media.

The capacity, price, and performance of optical storage continually improve, and IBM remains committed to providing its customers with these improvements over time. Even as new devices are introduced, the basic methods of accessing optical information remain consistent, as these new storage devices are being added under the current file system interfaces that optical storage programs have used for years.

Optical Media Types

Five categories of optical media are available to meet most storage requirements: CD-ROM, DVD-ROM, DVD-RAM, WORM (write-once read-many) optical cartridges, and erasable optical cartridges.

CD-ROM is a read-only format that is optimized for read performance. CD-ROMs are ideal for wide-scale distribution of programs and data. The CD-ROM data format is identical to the one that is used with personal computers. This makes it possible to develop CD-ROMs for use in both personal computers and the iSeries server. You can read CD-ROMs in either a CD-ROM or DVD drive.

DVD-ROM is a read-only format that provides a higher capacity than CD-ROM. Like CD-ROM, DVD-ROMs are excellent for wide-scale distribution of programs and data. You can only read DVD-ROMs in a DVD drive.

DVD-RAM is writable optical media that's available in both double-sided (Type I) and single-sided (Type II) formats, ranging from 2.6GB per cartridge to 9.4GB per cartridge. Both types can be accessed in a DVD-RAM drive and Type II media can be read in a DVD-ROM drive when the media is removed from the cartridge.

WORM storage is an economical way to archive data, yet still have it quickly and easily accessible. WORM media is available in 1x (650 MB), 2x (1.3 GB), 4x (2.6 GB), 8x (5.2 GB), and 14x (9.1 GB or 8.6 GB) capacities.

An erasable cartridge offers the most flexibility with similar capabilities as magnetic storage. Erasable media is available in 1x (650 MB), 2x (1.3 GB), 4x (2.6 GB), 8x (5.2 GB), and 14x (9.1 GB or 8.6 GB) capacities.

Both WORM and erasable cartridges must have a sector size of 1024 bytes per sector for 1x, 2x, and 4x media. For 8x media, permanent WORM must have a sector size of 2048. CCW (continuous composite write-once) WORM and erasable media can either be 1024 or 2048 bytes per sector. 14x media is available in CCW and erasable in 2048 or 4096 bytes per sector.

Optical System Configurations

All iSeries servers include a rack-mounted CD-ROM or DVD-ROM drive that is ideal for program and data distribution. The CD-ROM or DVD-ROM drive is primarily intended as a program and data delivery device and, even though many users can potentially access it simultaneously, it can access only one media at a time.

A DVD-RAM drive is available to replace the existing read-only drive or to be added as another internal drive. This drive can read CD-ROM or DVD-ROM media, and can read or write DVD-RAM media. The CD and DVD drives supported on iSeries cannot be used to write to CD-R, CD-RW, DVD-R, DVD-RW, or DVD+RW media.

An **optical media library** is a device that contains at least one optical disk drive and may contain many optical cartridges. Optical media libraries can manage large numbers of optical cartridges and users. "Supported Hardware for Optical Storage" on page 4 contains a list of all currently available optical media libraries that are supported by an iSeries server. These libraries support multiple numbers of users who access multiple numbers of optical cartridges simultaneously.

You can connect optical media libraries to your iSeries server in two ways: Directly-attached and LAN-attached.

Directly-Attached Libraries

One method of connecting optical media libraries is to directly connect the optical media library to your iSeries server. A multi-wire cable connects the library to an I/O processor or I/O adaptor card.

Directly-attached libraries support all of the following functions:

- All Hierarchical File System (HFS) application programming interfaces.
- Most integrated file system commands.
- Many OS/400 save and restore commands.
- Accessible by other LAN-connected iSeries servers by using the integrated file system (as shown in Figure 1 on page 3).




Figure 1. Example of Using an iSeries server Directly-Attached Library through a Local Area Network

LAN-Attached Libraries

The second method of connecting optical media libraries to your iSeries server is through a local-area network (LAN), as shown in Figure 2 on page 4.



Figure 2. Example of Using iSeries servers with a LAN-Attached Optical Library

LAN-connected optical media libraries can be used only through the HFS interface. The libraries can be accessed simultaneously either by several servers or any other devices on the LAN. LAN-connected optical media libraries have a controlling PC and do not require an iSeries server that acts as a controller. LAN-attached optical media libraries are often referred to as **optical servers**.

Supported Hardware for Optical Storage

A variety of hardware configurations for CD-ROM, DVD-ROM, and DVD-RAM are supported on iSeries servers. The table below lists the stand-alone optical devices available as well as the media supported in each.

Table 1. Supported stand-alone optical devices

Device Type	Hardware Resource Type and Model	Device	Media Supported
6320/6321	6320-002/6321-002	CD-ROM	CD-ROM CD-R
7210-020	6321-002	CD-ROM Bridgebox External device	CD-ROM CD-R

Table 1. Supported stand-alone optical devices (continued)

Device Type	Hardware Resource Type and Model	Device	Media Supported
6330 HH DVD-RAM	6330-002	DVD-RAM	CD-ROM CD-R ¹ CD-RW ^{1,2} DVD-ROM DVD-RAM ⁴
7210-025	6330-002	DVD-RAM Bridgebox External device	CD-ROM CD-R ¹ CD-RW ^{1,2} DVD-ROM DVD-RAM
6336	6336-002	DVD-ROM	CD-ROM CD-R CD-RW ² DVD-ROM DVD-R DVD-RAM ³
6331 Slim Rambo	6331-002	Slim Multi-Recorder	CD-ROM CD-R CD-RW ² DVD-ROM DVD-R DVD-RAM ⁵
6333 HH Rambo		IHH Multi-Recorder	CD-ROM CD-R CD-RW ² DVD-ROM DVD-R DVD-RAM ⁵
7210-030	6333-002	External device	CD-ROM CD-R CD-RW ² DVD-ROM DVD-R DVD-RAM ³
7212-102 Sam Bass	6330 6333 6336		CD-ROM CD-R CD-RW ² DVD-ROM DVD-R DVD-RAM ³
6337 Slim Line DVD-ROM		DVD-ROM	CD-ROM CD-R CD-RW ² DVD-ROM DVD-R DVD-RAM ³

Table 1. Supported stand-alone optical devices (continued)

Device Type	Hardware Resource Type and Model	Device	Media Supported
Notes:			
1. Read support in DVD-RAM drive only.			
2. Read support available at Version 5 Release 2 and later.			
3. Read support for Type II DVD-RAM 4.7GB media with the media removed only.			
4. Prior to V5R3, writes media in cartridge only. V5R3 on, writes media in or removed from cartridge.			
5. Write support for Type II DVD-RAM 4.7GB media with the media removed only.			

Optical media libraries come in a variety of configurations that are designed around the different forms of media and different connection options. Optical media libraries range from the single cartridge stand-alone model through models capable of holding 638 optical cartridges and twelve disk drives. Optical media libraries may be directly connected to the iSeries server for best functionality and performance, or may be connected through a LAN to allow independent access by PCs or other iSeries servers. Table 2 provides a list of supported optical library devices:

Table 2. Currently Supported Optical Storage Devices

Model	Drive Type	Connection	Cartridge Capacity	Number of Drives
3431-705	Multi-Function	LAN	1	1
3995-A23	Multi-Function	LAN	16	1
3995-022	WORM	LAN	32	2
3995-023	Multi-Function	LAN	32	2
3995-122	WORM	LAN	144	4
3995-123	Multi-Function	LAN	144	4
3995-C20	Multi-Function	LAN	20	1 or 2
3995-C22	Multi-Function	LAN	52	2
3995-C24	Multi-Function	LAN	104	2 or 4
3995-C26	Multi-Function	LAN	156	4 or 6
3995-C28	Multi-Function	LAN	258	4 or 6
3995-C40	Multi-Function	Direct	20	1 or 2
3995-C42	Multi-Function	Direct	52	2
3995-C44	Multi-Function	Direct	104	2 or 4
3995-C46	Multi-Function	Direct	156	4 or 6
3995-C48	Multi-Function	Direct	258	4 or 6
399F-100	Multi-Function	Direct	24-80	1-4
399F-200	Multi-Function	Direct	104-638	2-12

Note: For more information on connecting, configuring, and using device type 399F, see the Optical Storage web site at <http://www-1.ibm.com/servers/eserver/series/optical/>

The following table lists all the Input/Output attachment adapters that you can use to attach Direct and LAN optical libraries to your iSeries. You will need to verify which adapter is appropriate for your model iSeries server. For a list of the current supported attachments check the attachment summary link at our Internet Optical Storage home page:

<http://www-1.ibm.com/servers/eserver/iseries/optical/>

Table 3. Input/Output attachment adapters for optical library devices

Library Connection Type	iSeries Feature Code	Description/Comments
Token Ring	2724	4/16Mbps
Ethernet	2723/4723	10Mbps
Ethernet	2838/4838	100/10Mbps
Direct	2621	No longer supported
Direct	2729	Supported
Direct	2749	Supported
Direct	5702	Supported
Direct	6534	Supported

Optical Storage Organization

Optical storage is organized in three hierarchical units: volumes, directories, and files.

Optical Volumes

All optical data is stored on a unit that is called a **volume**. This is true regardless of the type of media, the type of optical media library that is used, and the way the storage device connects to your system. A single CD-ROM or DVD-ROM disk contains one volume. WORM and erasable optical cartridges are two sided and contain two volumes per cartridge (one per side). DVD-RAM can be either one sided or two sided.

Each volume has its own name that is chosen by the person who initializes the volume. The name chosen must be unique from the names of all other volumes on the system. Two volumes with the same name cannot be active at the same time. The volume name usually never changes after the volume is generated, although volume renaming is supported. The creator of the CD-ROMs and DVD-ROMs chooses their names, and the names cannot be changed.

HFS, the integrated file system, and the save and restore functions all use volume names to access or create data on the volume.

You can display and manage the volumes on an OS/400 by using the Work with Optical Volumes (WRKOPTVOL) command. The hierarchical and integrated file systems include the volume name in their path name to select which volume to use. A typical optical path looks like this:

```
/QOPT/VOLUMENAME/MYDIR/MYFILE.EXT
```

In the example:

- /QOPT is the name of the optical file system.
- /VOLUMENAME is the volume name that is chosen for the volume.
- /MYDIR is the name of a directory on the volume.

- /MYFILE.EXT is the name of a file in the directory.

Optical Directories

Information on an optical volume is organized into units called **directories**. A directory is a logical partition that can contain files and other directories called **subdirectories**. Every volume contains at least one directory called the **root** directory.

You are not required to have other directories below the root directory. Directories are a convenient way to keep files organized.

Optical Files

The basic element of optical storage is the optical **file**. Any data that is available to application programs can be stored in or retrieved from optical files in the form of a data stream. Optical files have the following features:

- Data is stored in a stream-file format.
- Data is independent of format and record structures.
- Data is accessed through byte offsets and lengths.
- Data is recognized and managed by the application that creates the file.

Using and Managing Optical Storage

You can display the primary menu for optical support by entering GO OPTICAL on the OS/400 command line. System administrators and programmers can access most optical commands through this menu. It is also convenient to enter many of the optical commands directly on the command line. These commands offer the following functions:

- Display optical volumes in a Directly-attached or LAN-attached optical media library device, CD device, or DVD device.
- Display files and directories that are contained in any directory in any optical volume.
- Display the file attributes of any optical file.
- Import or export media in a directly-attached optical media library, CD-ROM device, or DVD device.
- Make backup copies of a volume, directories, or files that are contained in a directly-attached optical devices.
- Initialize a volume that is contained in a DVD-RAM drive or in a directly-attached optical media library.
- Work with devices that represent optical media libraries, optical servers, CD drives, and DVD drives.
- Add, remove, or check the status of any LAN-attached optical server.
- Display active LAN-attached server conversations.
- Duplicate one optical volume to another
- Copy files and directories from one optical volume to another
- Check a volume for damaged directories and files

When you enter GO CMDOPT on the command line, a complete list of optical commands appears. Many of these commands are accessible through the previous GO OPTICAL menu.

Chapter 2. CD-ROM and DVD on iSeries Servers

This chapter provides information on configuring your CD-ROM, DVD-ROM, or DVD-RAM device as well as tips for loading and unloading media.

iSeries servers come with a rack-mounted CD-ROM or DVD-ROM drive. As an option, you can order a DVD-RAM drive as a feature to replace your internal drive or to have it in addition to your internal drive. All optical drives are multi-user devices that multiple users can access concurrently. The following interfaces are available to access the data on the CD and DVD devices:

- Save and restore interface
- The hierarchical file system application programming interface (HFS API)
- The integrated file system interfaces
- Optical commands and utility displays

Note: The CD-ROM and DVD-ROM drives on the iSeries server are not enabled for the digital audio disk format.

Supported Media Formats

OS/400 supports the reading of all CD and DVD media formatted with ISO9660 using the primary volume descriptor (PVD). OS/400 does not support ISO9660 extensions that use the supplementary volume descriptor (SVD). This level of support represents Level 1 compliance with the ISO9660 architecture (for a "Receiving System").

In addition, there is some support for ISO9660 extensions defined as the IEEE P1281 (System Use Sharing Protocol) and P1282 (Rock Ridge Interchange Protocol) specifications. These extensions are also known as Rock Ridge. OS/400 supports the Rock Ridge alternate name structures through the PVD. This support provides recognition and handling of mixed case and long file names, similar to UNIX[®]. Other structures defined by Rock Ridge are not supported by OS/400. To enable support for the Rock Ridge alternate name structures, use the Change Optical Attributes (CHGOPTA) command specifying EXTMEDFMT (*YES).

OS/400 also supports the reading of CD and DVD media formatted with Universal Disk (UDF) media format. UDF is the Optical Storage Technology Association (OSTA) supported subset of ISO/IEC 13346 and ECMA 167. OS/400 also supports writing to DVD-RAM media that is written with the UDF media format. See Chapter 10, "Optical Media Formats," on page 113 for more information about the different media formats.

Note: Although not accessible through OS/400, CD and DVD media that are created using the supplementary volume descriptor of ISO9660 may be accessible in an iSeries server. You can access them through an Integrated xSeries[®] Server for OS/400[®] server by running an operating system other than OS/400.

Configuring Your CD-ROM or DVD Drive

Depending on the model of iSeries server, you can position the CD-ROM or DVD drive either horizontally or vertically in the system.

Before you use the CD-ROM or DVD drive, you must have a device description for it. The system can create the device description automatically during an IPL if auto-configuration is on. Alternatively, you create it manually by using the Create Device Description Optical (CRTDEVOPT) command. Either method is acceptable. Once you create the device description, you can vary the configuration by using the Vary Configuration (VRYCFG) command. The configuration description for the CD-ROM or DVD device is *OPT. When the device description is varied on, it displays a status of ACTIVE.

Working with Volume Identifiers

Loading a CD-ROM or DVD media into a drive causes automatic reading of the information from the media. Part of this information is the **volume identifier**. The volume identifier is a name that was given to the CD-ROM or DVD-ROM when it was mastered. The identifier is also the name that is given to the DVD-RAM media when it was initialized. Depending on the media format, the volume identifier can be up to 32 characters in length. On the iSeries server, applications accessing data from the CD-ROM or DVD often refer to it by its volume identifier. For example, a CD-ROM volume identifier might be:

```
VOLID01
```

Applications that need to access file data from any optical media need to refer to the volume identifier. For example, you could write a C program to use the integrated file system APIs to read file /DIR1/FILE on the optical volume VOLID01. In this case, the application would specify path /QOPT/VOLID01/DIR1/FILE on the open request. For more information about application programming interfaces, refer to Chapter 7, "Using Application Programming Interfaces for Optical Support," on page 59.

Loading and Unloading CD-ROM and DVD Media

Loading a CD-ROM or DVD media into a drive causes reading of the media to retrieve the volume identifier. The system stores this volume identifier in an optical index database file to expedite future access to the media. Once the volume identifier is in the optical index, you can access the media through save and restore, as well as the application programming interfaces. Unloading the media from the drive removes the volume identifier from the optical index. The system adds the volume identifier to the optical index database in one of two ways:

- When the CD-ROM or DVD is loaded into a varied-on device. A CD-ROM or DVD can be loaded into a varied-off device. However, the optical index does not update until the device description gets varied on.
- When the user varies on a CD-ROM or DVD device description with media in it.

There is additional processing that could take several seconds after the tray slides in before the CD-ROM or DVD is usable. This is true for varying on as well. Even if the vary-on operation completes successfully, the CD-ROM or DVD is not usable until the system reads and stores the media information in the optical index. This may take several seconds after the vary-on operation completes.

When the user successfully loads a CD-ROM or DVD into a drive, the system sends the following message to the QSYSOPR message queue:

```
Volume VOLID01 added to optical device.
```

When a CD-ROM or DVD is successfully removed from a drive, the system sends the following message to the QSYSOPR message queue:

Volume VOLID01 removed from optical device.

You can use the Work with Optical Volumes (WRKOPTVOL) command to verify the successful adding or removing of the CD-ROM or DVD. You can also use the Display Optical (DSPOPT) command to display optical volume information.

Occasionally media may fail to load successfully into the drive. Listed below are the most likely causes for an unsuccessful load.

- Media or drive error occurred.
- The media format is not supported (digital audio CD-ROM).
- The system encountered a duplicate volume identifier.

Depending on the error, the tray may or may not eject if a CD-ROM or DVD fails to load. A failure to load the CD-ROM or DVD might not be obvious. Your first indication may be that you received one of the following messages when trying to access the CD-ROM:

```
Optical volume format not recognized.  
Optical volume not found.
```

If an error does occur when loading media or varying on the drive, the system signals an error message to the QSYSOPR message queue. This message describes the reason for the failure.

The processing for the load, unload, and vary-on operations runs in job QJOBSCD. If errors occur during these operations, view the job log of QJOBSCD to see the detailed messages.

Duplicate Volume Identifiers

With the availability of multiple CD-ROM and DVD drives on a single iSeries server comes the possibility of volume identifier collisions. Optical media, whether it is a CD-ROM, DVD, erasable, WORM, or virtual, are stamped or formatted with a volume identifier. The optical file system requires that all volume identifiers in an OS/400 be unique, due to the hierarchical directory structure of the QOPT file system. If a volume identifier collision occurs, the media containing the duplicate volume identifier becomes rejected with the following message that is signalled to QSYSOPR message queue.

```
Optical volume VOLID01 already exists.
```

In addition, the CD-ROM or DVD tray will eject the media.

Allocating the Device Description

The process of loading the CD-ROM or DVD media requires *SHRUPD use of the device description. Therefore, QJOBSCD must be able to obtain a *SHRUPD lock on the device description for the load to complete successfully. If another job is holding a conflicting lock on the device description, the load processing will fail with the following errors in the QJOBSCD job log.

```
Optical device xxxxx in use.  
Add optical disk cartridge failed to complete successfully.
```

As an example, assume that some job allocates OPT01 with an Exclusive Allow Read lock as follows:

```
ALCOBJ OBJ((OPT01 *DEVD *EXCLRD))
```

As long as the system holds this lock, CD-ROM and DVD loads will fail in QJOBSCD.

Non OS/400 access of optical devices

The iSeries optical drives are multi-server as well as multi-user devices. In addition to OS/400 jobs accessing the CD and DVD drives, non-OS/400 users running on another server can use these devices as well. For example, a Windows® server running on an Integrated xSeries Server can use the iSeries optical drive just as it does a local optical drive. The iSeries optical drive appears as a normal local optical drive in **My Computer** on Windows server. The use of a given drive may be exclusive to one server (or operating system) at a time depending on how the device is being accessed. In general, the drive can be shared concurrently between servers when both are only reading. If one server is writing to the drive, exclusive use is required. When a server has exclusive use of a drive, it actually has exclusive use of the optical volume in the drive, which can be determined through the Display Optical (DSPOPT) CL command. Enter the following to see which server has exclusive use of the volume in a specific optical drive:

```
DSPOPT VOL(*MOUNTED) DEV(OPTxx)
```

Page down to the second screen of Display Optical Volume Attributes. The *Volume held by* field indicates which server has exclusive use of the volume in the specified drive. If one server has exclusive use of the volume, another server or operating system cannot use the volume until it is released by the first server. If OS/400 attempts to use the volume while it is in use by another server, message OPT1790 (Operation not allowed or conflicts with another request) will be signalled. If a Windows server attempts to use the volume while in use by OS/400, an error will be signalled to the Windows user.

For more information on using the optical drive from a Windows server, see the Sharing devices topic in the Information Center.

For information on using the optical drive from a Linux partition, see the Linux topic in the Information Center.

Virtual Optical Devices

Beginning in Version 5 Release 2, a virtual optical device can be created with virtual media. Virtual media are real media that have been copied (or loaded) into an OS/400 stream file on internal disk storage and accessed, transparent to the application, through the virtual device. Virtual optical devices can be created in every OS/400 logical partition and used as an alternative to real devices to provide flexibility, better performance, and more functionality.

Beginning in Version 5 Release 3, writable virtual media can be created in virtual optical devices. Once created, loaded, mounted, and initialized, a virtual media can be accessed similar to real DVD-RAM media. Since these virtual media images exist as stream files in the integrated file system, they can easily be transported to other iSeries servers to be loaded and accessed. In addition, these media images can be copied to DVD-RAM media using the Duplicate Optical (DUPOPT) command or used to create CD-R media by sending the stream file to a PC and written using a CD burner. Another enhancement in Version 5 Release 3 is the capability to create up to 35 virtual devices in a single partition. This allows a hosting partition to support a virtual device for multiple other partitions concurrently.

| Below is an example of how to create a new writable virtual media using the
| Image Catalog commands:

| Create a virtual device description.
| CRTDEVOPT DEVD(OPTVIRT1) RSRNAME(*VRT)

| Vary on the virtual device description.
| VRYCFG CFGOBJ(OPTVIRT1) CFGTYPE(*DEV) STATUS(*ON)

| Create an integrated file system directory to store the image(s).
| CRTDIR DIR('/YourDirectory')

| Create an image catalog.
| CRTIMGCLG IMGCLG(TEST) DIR('/YourDirectory')

| Add a new image catalog entry (virtual media image) to the image catalog.
| ADDIMGCLGE IMGCLG(TEST) FROMFILE(*NEW) TOFILE(YourMediaImage)

| Load the image catalog on the virtual device.
| LODIMGCLG IMGCLG(TEST) DEV(OPTVIRT1)

| Initialize the virtual media.
| INZOPT NEWVOL(VOLUME1) DEV(OPTVIRT1)

| Virtual optical media VOLUME1 can now be used similar to a real DVD-RAM.
| The media image exists in file /YourDirectory/YourMediaImage
| WRKOPTVOL DEV(OPTVIRT1) CSI(*NONE)

| For more information about creating, using, and installing from virtual optical
| devices see the Virtual Optical topic of the Information Center.

Mastering Your Own CD-ROM

See the Optical Storage page (<http://www-1.ibm.com/servers/eserver/iseries/optical/cdrom/cddist.htm>) for instructions on CD mastering for iSeries.

Chapter 3. Directly-Attached Optical Media Libraries

This chapter provides information on configuring optical media libraries, labeling optical cartridges, and setting write protection on and off. This chapter also includes scenarios for working with optical media libraries.

Directly-attached optical media libraries are optical media libraries that are attached to the iSeries system through the SCSI** interface. On OS/400, multiple users can access data on directly-attached libraries concurrently. The following interfaces are available to access the data on these devices:

- Save and restore interface
- The hierarchical file system application programming interface (HFS/API)
- The integrated file system interfaces
- Optical commands and utility displays

Connectivity of non-IBM library devices

In addition to IBM 3995 optical libraries, you can now attach some non-IBM optical library devices to iSeries. For more information on which devices are supported and system configuration requirements, see the Optical Storage web site at <http://www-1.ibm.com/servers/eserver/iseries/optical/>

Configuring Optical Media Libraries

To create a device description for an optical media library device, use the Create Device Description (Media Library) (CRTDEVMLB) command. Specify device class *OPT. For example:

```
CRTDEVMLB DEVD(OPTMLB01) DEVCLS(*OPT) RSRNAME(OPTMLB01)
```

The configuration description for an optical media library device is *OPTMLB.

In addition, you can use the following commands to work with device descriptions:

- To change the device description, use the Change Device Description (Media Library) (CHGDEVMLB) command.
- To vary the device description, use the Vary Configuration (VRYCFG) command.
- To delete the device description, use the Delete Device Description (DLTDEVD) command.
- To work with the configuration status, use the Work with Configuration Status (WRKCFGSTS) command.

Labeling Optical Cartridges (WORM or Erasable)

Each optical disk cartridge contains two sides. Each side corresponds to an optical volume. There are two techniques to associate a volume ID label with the correct side of the optical cartridge. This is important to know when you set the write protect switch.

The optical cartridge should be labeled the first time the cartridge is added to an optical media library. This prevents any confusion in the future when you are attempting to determine which volume goes with which side.

Labeling a New Optical Cartridge with Uninitialized Volumes

To label a cartridge that is new (both volumes are uninitialized), follow these steps. The options are selected from the Work with Optical Volumes display.

1. Label sides A and B with the volume names you will use when initializing them.
2. Place the cartridge in the input/output station with side A facing up.
3. Add the optical cartridge by choosing option 1 (Add).

After the cartridge is added, the volume names appear on the Work with Optical Volumes display (see Figure 3) as system-generated IDs that consist of the system date and time.

```

Work with Optical Volumes
System: XXXXXXXX
Device . . . . . *ALL
Side information . . . . . : *ALL
Type options, press Enter.
  1=Add  2=Change  3=Copy  4=Remove  5=Display  8=Work with directories
 10=Initialize 11=Work with object links 12=Duplicate ...

Opt  Volume      Device      Volume      Media      Authorization
   -  -----      -          Type        Type        List
   -  941215111729 OPTMLB01    *UNFORMATTED *ERASE      QOPTSEC
   -  941215111730 OPTMLB01    *UNFORMATTED *ERASE      QOPTSEC

Parameters or command
====>
F3=Exit  F4=Prompt  F5=Refresh  F6=Print list  F9=Retrieve  F11=View 2
F12=Cancel  F14=Show extended information  F24=More keys ...
Bottom

```

Figure 3. Work with Optical Volumes Display

The earlier time corresponds to the side that was facing up in the input/output station. Therefore, if side A was facing up in the input/output station, then side A will have the earlier date and time of the two volumes.

4. Initialize the volume that corresponds to side A by choosing option 10 (Initialize) next to the system-generated volume ID. Use the labeled name. Repeat this step for side B.

Labeling an Optical Cartridge with an Initialized Volume

To label a cartridge that has at least one initialized volume on it, follow these steps. The options are selected from the Work with Optical Volumes display.

1. If the cartridge resides in an optical media library, remove it by choosing option 4 (Remove) next to the volume ID.
2. After the cartridge is removed, set one side of the cartridge to write-protected and the other side to write-enable. See “Setting Write Protection” on page 17 for more details on setting write protection.
3. Add the cartridge to an optical media library by choosing option 1 (Add).
4. Press F11 (View 2) to see the write-protected status of the newly added volumes.

5. Determine which volume is write-protected and make a record of this volume ID.
6. Remove the optical cartridge by choosing option 4 (Remove) next to the volume ID.
7. Label the write-protected side of the cartridge to the volume ID you previously recorded.

Setting Write Protection

The write-protect function prevents writing on the disk. A write-protect window shows when write protection is either on or off. To use the write-protect function, do the following:

1. Locate the write-protect switch on the cartridge.
2. Set the disk to read/write or read-only.
 - To make the disk read/write, move the write-protect switch to the off position. The write-protect window will be closed. You can write data on the disk.
 - To make the disk read-only, move the write-protect switch to the on position. The write-protect window will be open and data cannot be written on the disk.

A separate Write-Protect switch exists for each volume (side) on the cartridge.

Allocate/Deallocate Optical Drive

| Optical media libraries range from a model that has a single drive to a model that
| has twelve drives. Use the Work with Media Library Status (WRKMLBSTS)
| command to see the allocation status of each drive within a media library. The
| allocation status for a drive can be ALLOCATED, DEALLOCATED or
| *UNKNOWN.

The default allocation status for an optical drive is ALLOCATED, which means that the drive is available for use by the optical media library. The system sets this default value at IPL time. You can only change it using the Work with Media Library Status (WRKMLBSTS) command. DEALLOCATED means that the drive becomes unavailable for use by the optical media library. The allocation status for drives in a varied off optical media library is *UNKNOWN.

```

Work with Media Library Status
System: xxxxx

Type options, press Enter.
1=Vary on 2=Vary off 3=Reset resource 4=Allocate resource
5=Allocate unprotected 6=Deallocate resource 8=Work with description

Opt Device/ Resource Status Allocation Job
_ OPTMLB01 ACTIVE
6_ OPT03 OPERATIONAL ALLOCATED
_ OPT02 OPERATIONAL ALLOCATED

Bottom

Parameters or command
===>
F3=Exit F4=Prompt F5=Refresh F9=Retrieve F12=Cancel F17=Position to
F23=More options
(C) COPYRIGHT IBM CORP. 1980, 1999.

```

Figure 4. Work with Media Library Status

There are times when a drive should be removed from serving the optical media library such as when it is suspected of needing repair. To do this, you need to change the drive allocation status to DEALLOCATED. This will make the drive unavailable for use by the optical media library.

To deallocate a drive, select option 6 (Deallocate resource) on the desired device or resource.

```

Work with Media Library Status
System: xxxxx

Type options, press Enter.
1=Vary on 2=Vary off 3=Reset resource 4=Allocate resource
5=Allocate unprotected 6=Deallocate resource 8=Work with description

Opt Device/ Resource Status Allocation Job
_ OPTMLB01 ACTIVE
_ OPT03 OPERATIONAL DEALLOCATED
_ OPT02 OPERATIONAL ALLOCATED

Bottom

Parameters or command
===>
F3=Exit F4=Prompt F5=Refresh F9=Retrieve F12=Cancel F17=Position to
F23=More options
(C) COPYRIGHT IBM CORP. 1980, 1999.

```

Figure 5. Work with Media Library Status

The device allocation status of DEALLOCATED will remain until the drive is allocated again or the iSeries server is IPLed.

Getting Started with Optical Cartridges and Volumes

This topic introduces you to some of the optical support functions and familiarizes you with using optical cartridges and volumes. Through these examples you can learn how to:

- Work with optical volumes
- Add optical cartridges to an optical media library
- Initialize optical volumes
- Remove optical cartridges from an optical media library

Note: The device, volume, and directory names that are used in these examples are for illustrative purposes only. Your applications may require different volume names or different directory names. Additionally, your optical devices might have different names.

Only some of the available optical commands are discussed here. Do not use these topics as the primary reference for these commands because they do not describe all of the functions that are available. These topics provide a tutorial on getting started by using the optical utilities.

The following examples in this chapter assume that you have a new optical disk cartridge available to use and your optical media library is empty.

Adding Optical Cartridges to an Optical Media Library—Example

To add an optical cartridge to the optical media library you have attached to your iSeries server, place the optical cartridge with side A up in the input/output station of the optical library dataserver. Make sure the cartridge is seated properly. Use a new optical cartridge if one is available.

Note: It is not a requirement to have side A up to add an optical cartridge. To correctly label a new optical cartridge, however, it is important to know which side is up when you add the cartridge to an optical media library.

If the cartridge is new, you may want to label it before adding it for the first time. See “Labeling Optical Cartridges (WORM or Erasable)” on page 15 for more information.

Enter 1 (Add) in the options field and press the Enter key. The Add Optical Cartridge display appears. Complete the fields on the Add Optical Cartridge display as follows:

- Optical media library: Type the name of the optical media library to which you want to add the cartridge. This is the name of the library description that was created during installation. If you have more than one library attached, you need to know the optical media library association and configuration. To view your optical configurations, type `WRKMLBSTS MLB(*OPTMLB)`.

When all parameters are entered, press the Enter key.

Initializing Optical Volumes—Example

If the cartridge you added was a new cartridge, you will see something similar to the display shown in Figure 6. This indicates that the optical cartridge has two uninitialized volumes on it.

```

Work with Optical Volumes
System: XXXXXXXX
Device . . . . . *ALL
Side information . . . . . : *ALL
Type options, press Enter.
  1=Add  2=Change  3=Copy  4=Remove  5=Display  8=Work with directories
 10=Initialize 11=Work with object links 12=Duplicate ...

Opt  Volume      Device      Volume      Media      Authorization
   -  -----      -          Type        Type        List
   -  941215111729 OPTMLB01    *UNFORMATTED *ERASE      QOPTSEC
   -  941215111730 OPTMLB01    *UNFORMATTED *ERASE      QOPTSEC

Bottom

Parameters or command
====>
F3=Exit  F4=Prompt  F5=Refresh  F6=Print list  F9=Retrieve  F11=View 2
F12=Cancel  F14=Show extended information  F24=More keys ...

```

Figure 6. Work with Optical Volumes Display

An **uninitialized volume** is an optical volume that has never been formatted or initialized. It is similar to a new diskette that needs to be formatted. If a volume has not been initialized, it has a volume type of *UNFORMATTED (uninitialized).

An uninitialized volume does not have a volume name written to it. When an uninitialized volume is added to an optical media library, a volume name that consists of a date and time (YYMMDDHHMMSS) is assigned to it. Optical volumes cannot be written to or read from until they are initialized.

On the Work with Optical Volumes display, you see that two volumes were added, one for each side of the optical disk cartridge. If either of the volumes was initialized, the volume type would indicate *PRIMARY or *BACKUP, and most likely the volume ID would be different from the ones in the example. An optical disk can contain one volume that is initialized and one volume that is uninitialized. The volumes are treated independently even though they exist on the same cartridge.

If both volumes are uninitialized, the earlier date and time for the volume names indicates the volume that was facing up in the input/output station. In this example, because side A was facing up when the add operation was performed, we know that 941215111729 is side A. This is important to know so the volume can be properly labeled when removed from the library.

To initialize the volume you determined to be side A, enter option 10 (Initialize) in the options field beside that volume. Press the Enter key.

Attention: Initializing a previously initialized volume makes all existing data on that volume inaccessible. If you have entered 10 next to a volume that is already initialized and you do not want to lose the data on that volume, *do not* continue with this function. Use a volume that is uninitialized.

Complete the fields on the Initialize Optical Volume display as follows:

- Volume identifier: This is the existing volume ID of the volume that you are going to initialize. For uninitialized volumes, this name is a system-generated name that consists of the date and time. This is only a temporary name until the volume initializes.
- New volume identifier: Type the new name for the specified volume. This is the name that users and applications will use when referring to the volume.
- Volume-full threshold: Leave this set to the default value. Use this value to give the volume a logical volume-full threshold.
- Check for active volume: Leave this set to *YES to verify that the optical volume was previously initialized.
- Clear: Leave this set to *NO. This specifies whether or not existing data on the volume will be cleared during the initiating process. This parameter only applies when the volume media type is *DVD-RAM. Specifying CLEAR(*YES) can cause this operation to take up to one hour.
- Text description: Enter a short description of the volume. Fifty characters are available for this field.

Press the Enter key to initialize the volume. After the volume is initialized, the Work with Optical Volumes display reappears as shown in Figure 7.

Note: Initializing an erasable optical volume can take up to 30 minutes.

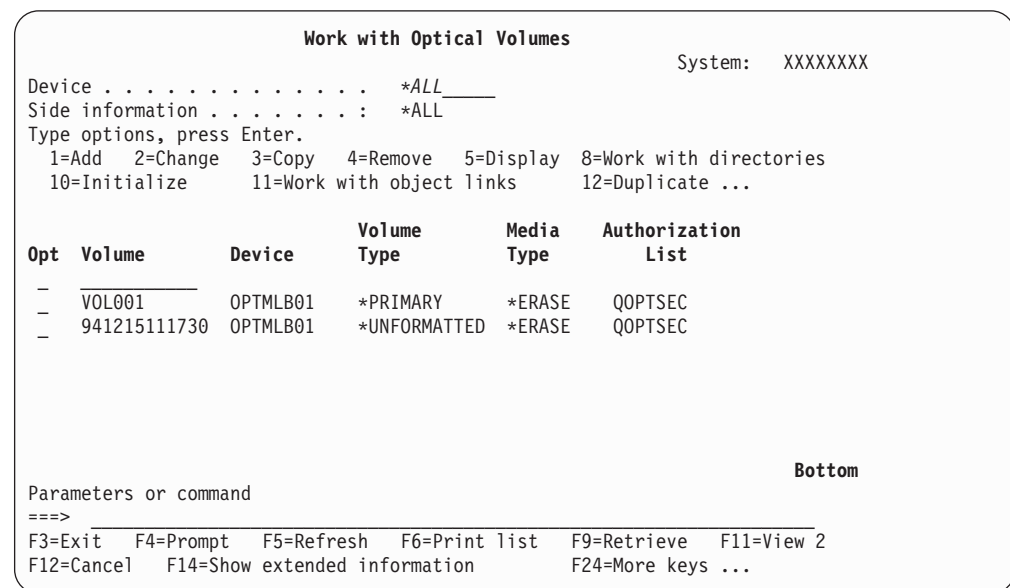


Figure 7. Work with Optical Volumes Display

The previously uninitialized volume is initialized and the type has changed from *UNFORMATTED to *PRIMARY. The volume is now available for reading and writing files and directories.

From this display, you can perform the following functions:

- Add an optical cartridge by selecting option 1 (ADD).
- Change the volume attributes by selecting option 2 (Change).
- Copy a volume by selecting option 3 (Copy).
- Remove an optical cartridge by selecting option 4 (Remove Cartridge).
- Display volume information by selecting option 5 (Display).

- Print volume information by selecting option 6 (Print).
- Rename a volume by selecting option 7 (Rename).
- Work with directories on the volume by selecting option 8 (Work with Directories).
- Delete volume information for a previously removed volume by selecting option 9 (Delete).
- Initialize or reinitialize a volume by selecting option 10 (Initialize).
- Work with directories and files on the volume by selecting option 11 (Work with object links).
- Duplicate a volume to another optical media by selecting option 12 (Duplicate).
- Check the volume for damaged files and return a count of the directories and files by selecting option 13 (Check).

Removing an Optical Disk Cartridge—Example

This example shows how to remove the optical disk cartridge from the optical media library. Because there are two volumes on a cartridge, removing a cartridge actually removes two volumes. Therefore, as shown in the example that follows in Figure 8, specifying option 4 (Remove) on either VOL001 or 941215111730 produces the same result.

```

Work with Optical Volumes
System: XXXXXXXX
Device . . . . . *ALL_____
Side information . . . . . : *ALL
Type options, press Enter.
 1=Add  2=Change  3=Copy  4=Remove  5=Display  8=Work with directories
10=Initialize 11=Work with object links 12=Duplicate ...

Opt  Volume      Device      Volume      Media      Authorization
   -  _____  _____  Type        Type        List
   -  VOL001      OPTMLB01   *PRIMARY    *ERASE      QOPTSEC
   -  941215111730 OPTMLB01   *UNFORMATTED *ERASE      QOPTSEC

Bottom

Parameters or command
====>
F3=Exit  F4=Prompt  F5=Refresh  F6=Print list  F9=Retrieve  F11=View 2
F12=Cancel  F14=Show extended information  F24=More keys ...

```

Figure 8. Work with Optical Volumes Display

Before removing a cartridge, make sure that the input/output station is empty. The optical media library cannot remove a cartridge if the input/output station is occupied by another cartridge.

Enter 4 (Remove) in the Opt field next to the cartridge you wish to remove and press the Enter key. The Remove Optical Cartridge display appears. Complete the fields of the Remove Optical Cartridge display as follows:

- Volume Identifier: This is the volume you selected on the Work with Optical Volume display.
- Volume description option: Specify *REMOVE for this field to remove the volume descriptions from the optical index database files after the cartridge has been

removed. Specify *KEEP for this field to save the volume descriptions for initialized volumes in the optical index database files. This causes the system to consider the volumes as *REMOVED.

Press the Enter key to remove the optical disk cartridge. The optical disk cartridge has now been moved to the input/output station of the optical library dataserer.

Press F3 (Exit) to return to the optical support main menu. The screen that displays shows that the optical disk cartridge is no longer available.

Press F3 again to return to the OS/400 command line.

Changing Optical Environment Parameters

The Change Optical Attributes (CHGOPTA) command can be used to change specific optical configuration parameters that affect all jobs using the optical file system. The following parameters can be set according to your desired configuration.

- Copy attributes

Specifies whether to copy the file attributes of the source file when copying or moving files between the QDLS and the QOPT file systems using the HFS copy or move stream file APIs. This parameter is ignored when copying or moving files within the optical file system and is not valid for LAN-attached optical libraries. This value should be set to *NO when applications do not require the file attributes to be maintained when copying between the QOPT and QDLS file systems. Using this option improves the performance of move and copy operations and reduces the optical storage requirements when writing to an optical disk.

- Held file attribute

Specifies whether held optical file support is enabled or disabled for the /QOPT file system. When held optical file support is disabled it is up to the user to ensure that correct recovery procedures are followed when there is a close file error condition. Refer to Appendix C, "Recovering Held Optical Files," on page 137 for a further explanation of held optical files.

- Allow variant characters

Indicates whether or not variant characters can be specified for path names when accessing files on optical volumes in Universal Disk Format (UDF) through the Hierarchical File System (HFS) interfaces. This parameter does not affect access to optical volumes formatted in High Performance Optical File System (HPOFS) or ISO9660 formats. It also has no affect on integrated file system interfaces for all optical media formats.

- Support extended media formats

Indicates which media format to use when a media is added to an optical device. Some media have both a primary and extended media format. For instance, a UDF-Bridge Disk has ISO9660 as the basic format and Universal Disk Format (UDF) as the extended format. Rock Ridge alternate name structures is another format that is supported. This parameter will let the user choose which format is to be used when a media is added to a device.

Chapter 4. LAN-Attached Optical Media Libraries

This chapter describes aspects of optical LAN support, including:

- Configuration of optical LAN support.
- Volume security.
- Description of LAN-specific commands.
- Differences between directly-attached and LAN-attached devices.
- Getting started with optical LAN support

Optical LAN support allows one or more iSeries server access to optical devices or libraries over either a token-ring or Ethernet LAN connection. Figure 9 on page 26 shows an example optical configuration where three LAN-attached optical media libraries are shared by two servers and a workstation.

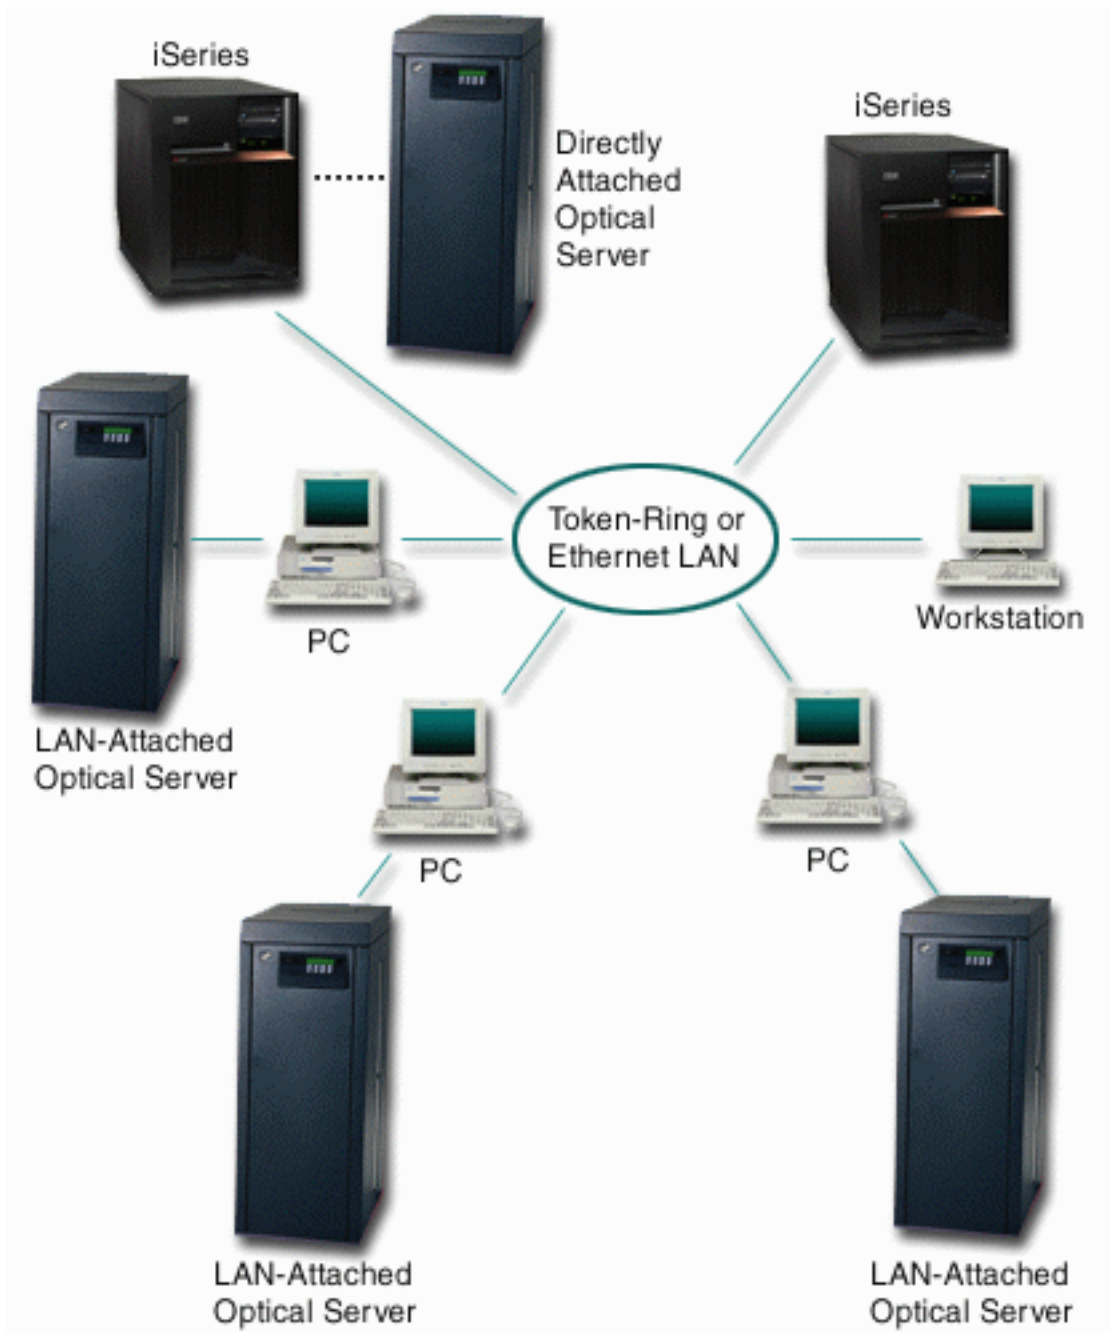


Figure 9. Typical Optical LAN Configuration

The following interfaces are available to access data on optical LAN devices:

- Hierarchical file system application program interfaces (HFS APIs)
- Optical commands and utility panels

Accessing data on a LAN device works the same way as accessing data on a directly-attached device. To access or store information, you need to know the path name. The path name consists of the file system name, volume, directory, and file to use. Optical support determines where to store the volume, and whether the device is directly-attached or LAN-attached. The HFS request gets routed to the appropriate server.

An application that uses HFS APIs to access optical volumes in directly-attached devices can also access volumes in remotely-attached optical LAN devices. Volume names must be unique across all optical devices, whether the volumes exist in a directly-attached device or a LAN device. Most HFS APIs that are supported by directly-attached devices are also supported by LAN devices with minor differences. For a complete list of what is supported, unsupported, and changed, refer to Table 5 on page 60.

Three commands are dedicated to LAN support. They are used for adding, removing, and displaying optical LAN devices.

- Add Optical Server (ADDOPTSVR) command
- Display Optical Server (DSPOPTSVR) command
- Remove Optical Server (RMVOPTSVR) command

There are other optical commands and utility panels that support optical LAN devices and volumes in those devices. For a complete list of these commands, see Table 4 on page 37.

Configuring Optical LAN Support

To start using optical LAN support, all of the hardware and software items must be in place and active. Figure 10 on page 28 shows the layers of software used in optical LAN support. All software required for the iSeries server is part of the base OS/400 operating system. OS/400 Extended Base Support should also be installed. Use the GO LICPGM menu to determine whether or not the extended base support is installed. All C2x series libraries come with an OS/2[®] PC loaded with all necessary software. For more details on LAN hardware, refer to “Supported Hardware for Optical Storage” on page 4.

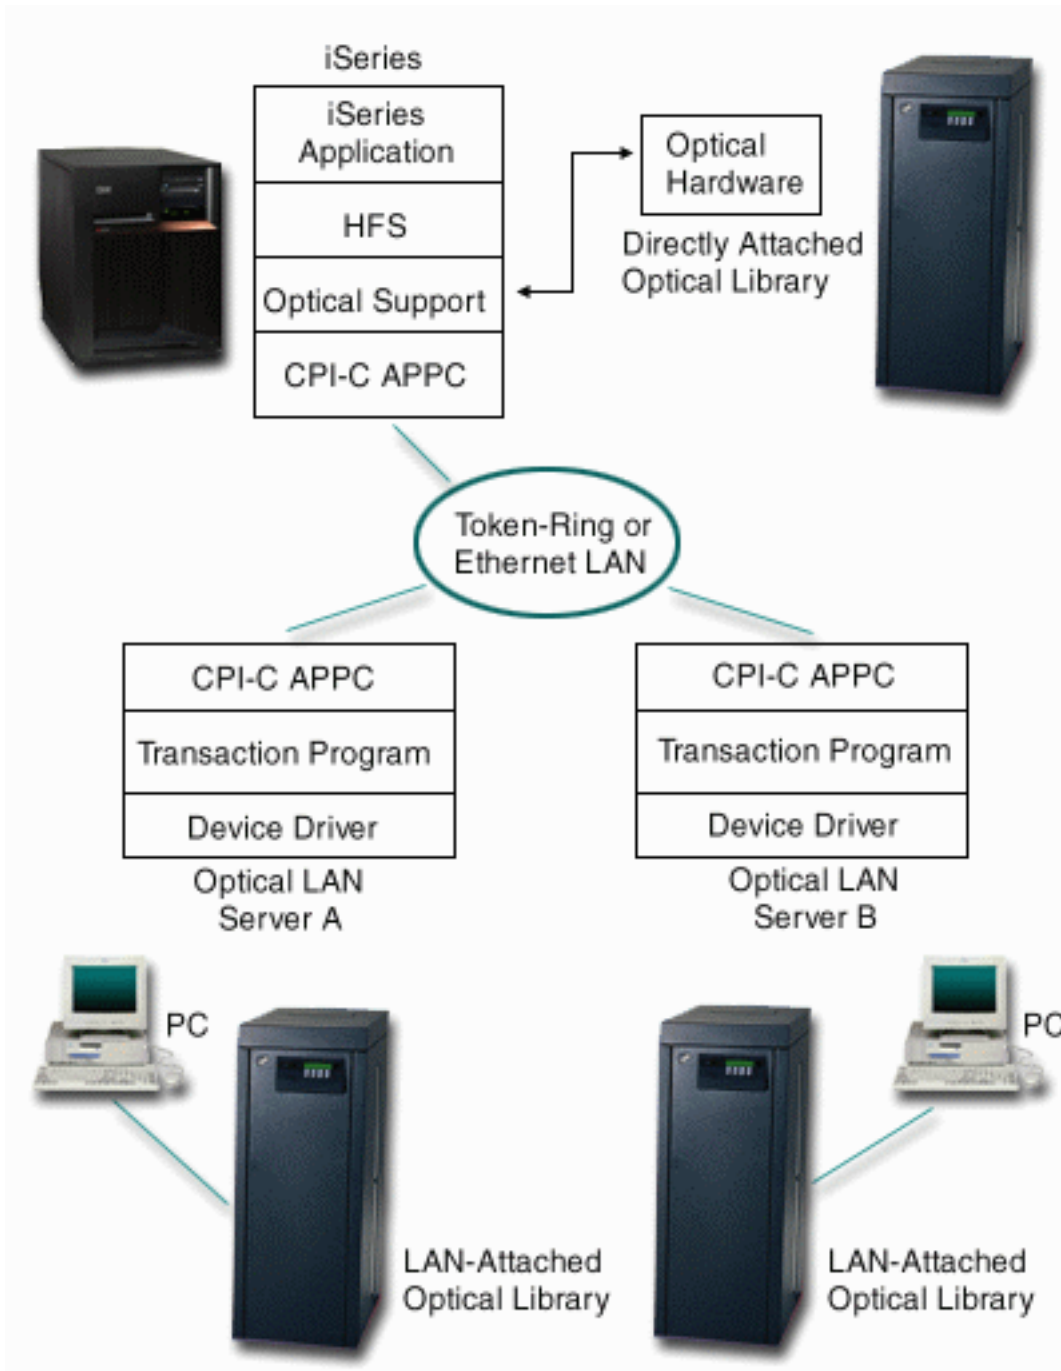


Figure 10. Optical LAN Software Interfaces

Advanced program-to-program communications (APPC) is used to communicate between an OS/400 and the optical LAN device. APPC can be run over a SNA network or over a TCP/IP network. On OS/400, the LAN device is configured as an APPC controller and device which is on either a token-ring or Ethernet connection. A sample SNA configuration to use when configuring a LAN device is given in “OS/400 Communications Control Setup—Example” on page 29. Figure 11 on page 29 shows the relationship between OS/400 communications control setup and the LAN device controller setup. See “Configuring Optical LAN Support over a TCP/IP connection” on page 30 for more information on TCP/IP configuration.

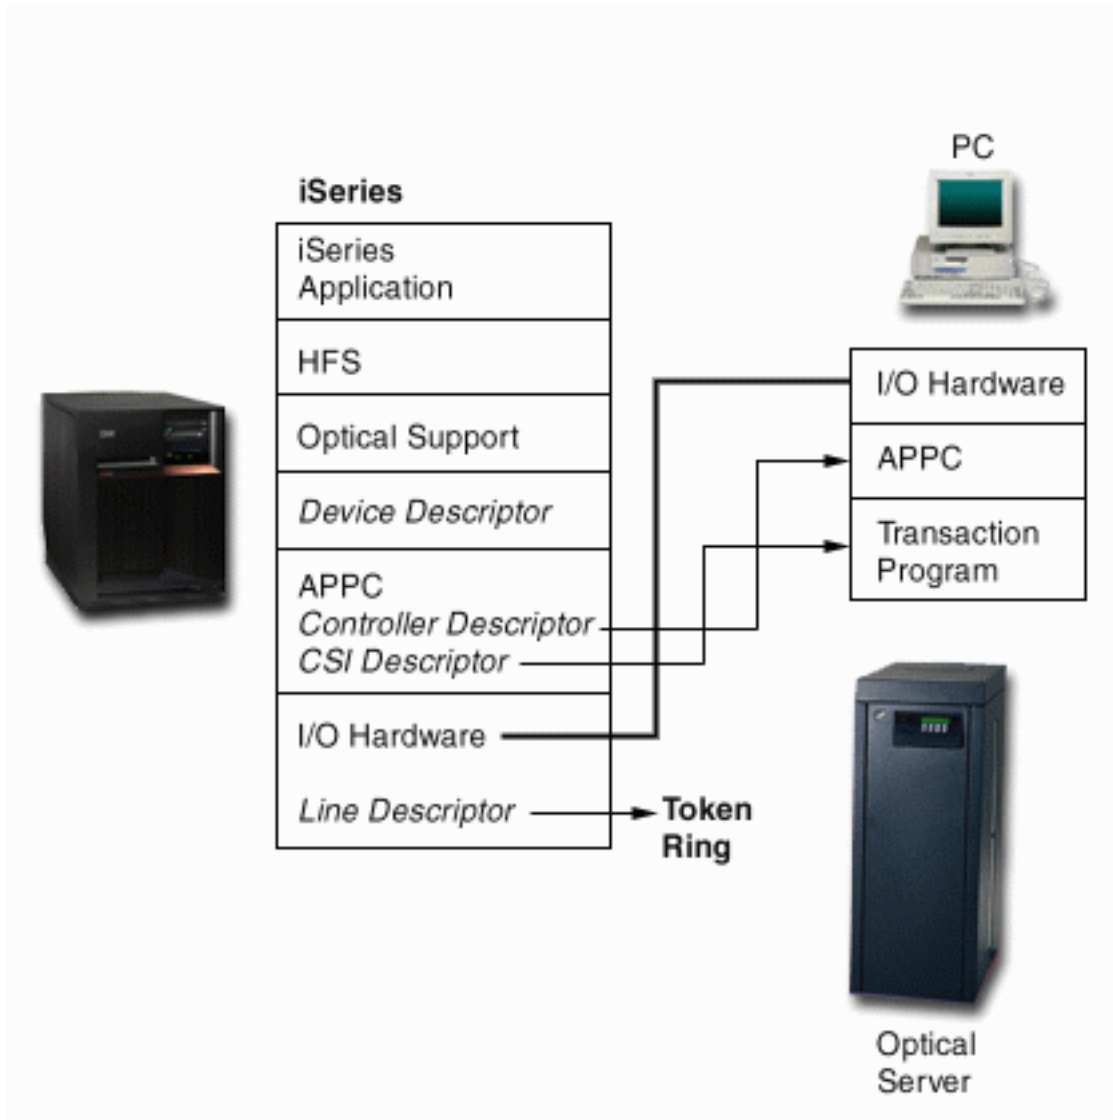


Figure 11. Optical LAN Configuration

OS/400 Communications Control Setup—Example

The following example shows you how to set up optical LAN support on an iSeries server. If you need more information on commands or parameters, see the Communications Configuration book.

1. Use the Create Line Description Token Ring (CRTLINTRN) command to define the token ring that is used to access the optical LAN device. Use the DSPHDWRSC TYPE(*CMN) command to find the resource name (LIN031 in this example) and a token-ring port. If you plan on using an existing line, you do not need to create a line description.

```
CRTLINTRN LIND(TRLANOPT)
RSRCNAME(LIN031)
LINESPEED(16M)
MAXFRAME(16393)
LINKSPEED(16M)
AUTOCRTCTL(*NO)
TEXT('Token Ring Line for Optical LAN device')
```

2. Use the Create Controller Description APPC (CRTCTLAPPC) command to create the controller description for the logical definition for the PC. The values of ADPTADR and RMTCPNAME are dependent on the PC configuration. Refer to the *IBM 3995 LAN Optical Library Dataserver: Reference* for more details.

```
CRTCTLAPPC CTLD(OPTSRV2)
LINKTYPE(*LAN)
APPN(*NO)
SWTLINLST(TRLANOPT)
RMTCPNAME(LU3995X)
ADPTADR(10005AE8B454)
TEXT('Controller for Optical LAN device')
```

3. Use the Create Device Description APPC (CRTDEVAPPC) command to create the device description for the logical definition for the optical LAN device.

```
CRTDEVAPPC DEVD(OPTSRV2)
RMTLOCNAME(LU3995X)
CTL(OPTSRV2)
APPN(*NO)
TEXT('Device Descriptor for Optical LAN device')
```

4. Use the Create Communications Side Information (CRTCSI) command to define the communications side information (CSI) object that logically defines the optical LAN device. The name of the CSI is referred to as the optical LAN servers destination name. The CSI also defines the name of the transaction program which is started at the LAN device. The value for TNSPGM is dependent on the configuration of the PC. Refer to the *IBM 3995 LAN Optical Library Dataserver Reference* for details. The library which contains the CSI definitions must be in the library list of any user who will use the optical LAN device. You also need *USE authority to the *CSI object before you can use the optical LAN device which it defines.

```
CRTCSI CSI(<user lib>/OPTSRV2)
RMTLOCNAME(LU3995X)
TNSPGM(HFSSRV)
DEV(OPTSRV2)
LCLLOCNAME(*LOC)
TEXT('CSI description of Optical LAN device')
```

5. After all the create commands have completed successfully, use the Work with Configuration Status (WRKCFGSTS) command to make sure that the line, controller, and device descriptions are varied on.
6. Use the Add Optical Server (ADDOPTSVR) command to enable LAN support for a maximum of 16 CPI Communications destinations.

```
ADDOPTSVR CSI(OPTSRV2)
```

LAN Controller Setup

The LAN controller, usually a PC, must be configured to allow iSeries server requests to communicate with the controller as an optical LAN device. For details on LAN controller configuration, see the *IBM 3995 LAN Optical Library Dataserver Reference*

Configuring Optical LAN Support over a TCP/IP connection

Optical LAN for iSeries support uses Advanced Program to Program Communications (APPC) to communicate over System Network Architecture (SNA). APPC is used on both OS/400 and OS/2. Optical LAN support does not currently support TCP/IP. If you want to use your Optical LAN library in a TCP/IP network, you can enable AnyNet[®] support on both the iSeries server and the OS/2 PC controller.

The AnyNet product on the iSeries server is called AnyNet/400. AnyNet/400 is shipped with the OS/400.

AnyNet products allow application programs written for one communications protocol to run over non-local protocols without changing (or recompiling) the application programs. The destination address determines if the request is sent over the local protocol or through the AnyNet code and on to a non-local protocol.

AnyNet/400 allows sockets, intersystem communications function (ICF), CPI Communications (CPI-C), and CICS/400[®] applications to run over APPC, TCP/IP, and Internetwork Packet eXchange (IPX). AnyNet/400 is based on the Multiprotocol Transport Network (MPTN) architecture, and is designed to allow any application to run over any networking protocol. You can use AnyNet to:

- Access APPC using TCP/IP if your applications were developed for system network architecture (SNA) but you are using TCP/IP to connect the systems.
- Access APPC using IPX if your applications were developed for SNA but you are using IPX to connect the systems.
- Access sockets using SNA if your sockets applications were developed for TCP/IP but you are using SNA to connect the systems.
- Access sockets using IPX if your sockets applications were developed for TCP/IP but you are using IPX to connect the systems.

An AnyNet TCP/IP configuration example for Optical LAN support can be found off of the LAN Configuration using TCP/IP link at our Internet Optical Storage home page or in the Information Center AnyNet topic.

Volume Security

You can use authorization lists to secure volumes within a LAN device in the same way that you secure volumes in a directly-attached optical device. Chapter 6, "Optical Security and Auditing," on page 53 describes this further. If multiple iSeries servers use an optical LAN device, you must define volume security on each iSeries server. Volume security for the first iSeries server can be exactly the same, slightly different, or totally different from the security that is defined on the second iSeries server. It is up to the security administrator of each iSeries server to define their security requirement for the optical volumes.

When an optical LAN device is initially added to the system using the Add Optical Server (ADDOPTSVR) command, all volumes in that device are secured by the default authorization list QOPTSEC. The default authorization list for a volume can be changed by using the Change Optical Volume (CHGOPTVOL) command.

Before removing an optical LAN device from the iSeries server using the Remove Optical Server (RMVOPTSVR) command, you should determine whether or not you want to retain the list of volumes along with their authorization lists in the optical index database. The VOLOPT parameter allows you to either keep or remove the volume entries for a LAN device from the optical index database. If you choose to have the entries removed, you need to redefine the authorization list for any volume not defined by the default authorization list the next time the server is added for use.

Description of LAN-Specific Commands

Three CL commands are dedicated to optical LAN support:

- Add Optical Server (ADDOPTSVR)
- Remove Optical Server (RMVOPTSVR)
- Display Optical Server (DSPOPTSVR)

Add Optical Server

The Add Optical Server (ADDOPTSVR) command enables optical support to access an optical LAN device. Before you issue the ADDOPTSVR command, make sure that the communications controller and device that represent the optical LAN device are varied on and are active. Typically, the command should be run during the following scenarios:

- After initial installation
- After an IPL
- To restart optical LAN support for a server that has been removed by using the RMVOPTSVR command

The ADDOPTSVR command performs the following functions:

- Verifies that the CPI Communications destinations that define optical LAN devices are correct. Currently, the limit on the number of optical LAN devices is 16.
- Sets up the index for the names of all the volumes in all optical LAN devices.
- This command verifies that all optical volume names are unique for both LAN-attached and directly-attached optical devices. The system issues diagnostic error messages for any duplicate volume names that are found. If there are any duplicate volume names, only the first volume with the duplicate name is accessible.

After the ADDOPTSVR command has completed successfully, optical LAN support knows all the added optical LAN devices. Dynamic checking of any processed HFS request determines if the request is for a volume in a LAN server or a directly-attached device.

Remove Optical Server

The Remove Optical Server (RMVOPTSVR) command disables access to optical volumes that are in a specific LAN device, a list of LAN devices, or all LAN devices. Access to volumes in directly-attached devices is not affected. The RMVOPTSVR command performs the following functions:

- This command ends any active conversations that belong to the current job that issues the command.
- This command verifies that there are no other allocated LAN conversations before attempting to end.
- When the last server has been removed, this command indicates that LAN support is no longer active.

Display Optical Server

The Display Optical Server (DSPOPTSVR) command can be used to determine if LAN support is active and to display information about the optical LAN configuration. The DSPOPTSVR command has two different options:

- *DEST (destinations). This option displays all the CPI Communications destinations that were added with the Add Optical Server (ADDOPTSVR) commands and their current status.
- *CONV (conversations). This option displays all active conversations, the destination of each conversation, the jobs using the conversations, and the path of any open files.

Differences between Directly-Attached and LAN-Attached Devices

There are some functional differences between optical LAN devices and directly-attached optical devices of which you should be aware. If you have both types of devices, please be sure to review the following list.

- You must install the OS/400 Extended Base Support before you can add an optical LAN device.
- LAN devices allow fewer HFS commands than directly-attached devices. Additionally, there are some differences between the level of support that the two attachments (directly-attached or LAN-attached) provide for the commands. Be sure to review Table 5 on page 60, which identifies these differences.
- LAN optical volumes do not support integrated file system APIs and commands.
- LAN optical volumes do not support save interfaces and restore interfaces.
- The term *path* refers to a file-system name, volume name, directory name, and file name. For directly-attached devices, the path cannot contain more than 294 characters. For optical LAN devices, the path cannot contain more than 256 characters.
- The valid character set for path names is slightly different between LAN and direct. For the character set that LAN supports, see the *IBM 3995 LAN Optical Library Dataserver: Reference*. For the character set that directly-attached libraries support see "Path Names" on page 87.
- Depending on the initialization of a LAN volume, path and file names may or may not be case-sensitive. Path and file names that are created on a volume in a directly-attached device are not case-sensitive. The path
/directory1/directory2/file1

or

/DIRECTORY1/DIRECTORY2/FILE1

refers to the same file when that file was created on a direct device. The path may or may not refer to the same file when that file was created on a volume that is located in a LAN device.

- A single iSeries server job using an optical LAN device may have a maximum of 249 open files for each optical LAN device.
- LAN does not support extended attributes for files, and directly-attached devices do.
- LAN support does not allow file attribute information on the open stream file request.
- LAN does not support all of the optical commands. For a list of supported optical commands and differences, refer to Table 4 on page 37.
- LAN does not support the expanding buffer I/O, the attribute that is named QOPT.IOMETH, on the open stream file request. Any specifying of the attribute remains ignored.

- The available volume space returns differently for LAN-attached devices than for directly attached devices. LAN support returns a 44-byte output buffer whose format is defined in “Retrieve Volume Information Function” on page 70.
- Support exists for using the HFS API Copy Stream File (QHFCPYSF) to copy between two optical LAN devices. Copying files between two directly-attached optical devices is supported. No support exists for copying files between a LAN device and a directly-attached device. To do this, you need one additional step. As an example, you could first copy the files from either the LAN device or directly-attached device to a folder file on the iSeries server. Appendix F, “Programming Examples,” on page 155 has an example tool for copying between file systems. You would then copy this folder from the iSeries server to the other optical device.
- Unlike direct support, HFS LAN does not receive information when a cartridge imports or exports from a dataser. The LAN is unaware of an imported cartridge that occurred since the last time the Add Optical Server (ADDOPTSVR) command was issued. This remains true only until the LAN receives a request for a volume on the cartridge. At that time, LAN automatically performs a query of all active optical LAN devices in order to find the volume. The LAN support receives no information when cartridge removal occurs from a LAN device. If a request for a volume on the cartridge occurs, the LAN device returns an error condition of Volume not found. If this is a problem, update the LAN support each time a cartridge is removed. Do this by either performing an Update LAN request UPD/LAN request through the QHFCTLFS program, or by issuing the Add Optical Server (ADDOPTSVR) command.

Getting Started with Installing an Optical LAN Device

The following topics illustrate some of the basic functions you may attempt when initially adding an optical LAN device to your system. The example illustrates the following:

- Adding an optical LAN device
- Working with volumes in a LAN device
- Working with directories
- Working with files
- Removing an optical LAN device

It is assumed that the configuration of the iSeries server and the optical controller device has been previously done and that the communications line, controller, and device are in an active state. It is also assumed that the communications side information (CSI) object is defined and you have been informed of its name.

Adding an Optical LAN Device

Figure 12 on page 35 shows an attempt to add optical server OPTSVR2 to a configuration using the Add Optical Server (ADDOPTSVR) command. The ADDOPTSVR command attempts to contact the communications device defined in the CSI object named OPTSVR2. If the command does not appear to complete successfully, start your problem analysis by checking for messages in the message queue for QSYSOPR. Also check for messages indicating that you have optical volumes with the same name. If duplicate volumes are found, only the first volume found can be accessed.

```

Add Optical Server (ADDOPTSVR)

Type choices, press Enter.

Side information . . . . . optsvr2_      Name
Library . . . . . *LIBL_____      Name, *LIBL, *CURLIB
      + for more values      _____
      *LIBL_____

Bottom
F1=Add  F2=Change  F3=Copy  F4=Remove  F5=Display  F8=Work with directories
F10=Initialize  F11=Work with object links  F12=Duplicate ...

```

Figure 12. Add Optical Server

Working with Volumes in a LAN Device

Wait for the ADDOPTSVR command to complete successfully. Then use the Work with Optical Volumes (WRKOPTVOL) command to view a list of volumes that is found in the optical device that is defined by the OPTSVR2 CSI object. To see a list of volumes in all directly-attached devices and all LAN-attached devices, issue the following command:

```
WRKOPTVOL DEV(*ALL) CSI(*ALL)
```

A display similar to the one shown in Figure 13 is shown.

```

Work with Optical Volumes
System:  XXXXXXXX
Device . . . . . *ALL_____
Side information . . . . . : *ALL
Type options, press Enter.
  1=Add  2=Change  3=Copy  4=Remove  5=Display  8=Work with directories
 10=Initialize  11=Work with object links  12=Duplicate ...

Opt  Volume      Device      Volume      Media      Authorization
-----
-  VOLL01      OPTSVR3    *SVRVOL      QOPTSEC
-  VOLL02      OPTSVR2    *SVRVOL      QOPTSEC
-  VOLL03      OPTSVR2    *SVRVOL      QOPTSEC
-  VOLL04      OPTSVR2    *SVRVOL      QOPTSEC
-  VOLL05      OPTSVR2    *SVRVOL      QOPTSEC
-  VOLL06      OPTSVR2    *SVRVOL      QOPTSEC
-  VOLUMEXXXX > OPTSVR2    *SVRVOL      QOPTSEC
More...

Parameters or command
====>
F3=Exit  F4=Prompt  F5=Refresh  F6=Print list  F9=Retrieve  F11=View 2
F12=Cancel  F14=Show extended information  F24=More keys

```

Figure 13. Work with Optical Volumes Display

In the figure, one of the volume names ends with a greater than sign (>). This signifies that the name of the volume is longer than what can be shown on the

display. To see the full name of the volume, place the cursor on the name and press function key 22 (F22). A pop-up window is displayed that shows the entire volume name.

The Work with Optical Volume display shows options that are for both directly-attached and LAN-attached devices and volumes. If you attempt an option that is not supported by a LAN device, you receive message OPT1346. This indicates that the operation is not allowed to a volume that is located in a remote optical device.

Media Interchange Between LAN and Direct Attach Libraries

The following interchange rules apply for volumes initialized in a directly attached optical library:

- WORM volumes are supported in a LAN attach library.
- Rewritable volumes initialized in HPOFS media format are supported in a LAN attach library.
- Rewritable volumes initialized in UDF media format are not supported in a LAN attached optical library.

The following interchange rules apply for volumes initialized in a LAN attached optical library:

- WORM volumes are supported in a direct attach library.
- Rewritable volumes initialized as WORM format and uppercase characters are supported in a direct attach library.
- Rewritable volumes initialized as rewritable format are not supported in a direct attach library.
- Rewritable volumes initialized to use mixed uppercase and lowercase characters are not supported in a direct attach library.

Removing an Optical LAN Device

If you no longer need access to a LAN-attached device and the volumes that the device contains, you can use the Remove Optical Server (RMVOPTSVR) command. You can use this command for a single CSI, a list of CSIs, or all CSIs. Additionally, you can keep or remove the record of all volumes that are found in devices that are represented by the CSIs from the iSeries server. Choosing to remove the volume entries from the iSeries server also removes any special security for those volumes. You must redefine the special security for the volumes when those volumes are again accessible by the iSeries server.

Chapter 5. Working with Optical Volumes

This chapter describes the "Work with Optical Volumes" options on the Optical Support Utilities main menu. These options are organized hierarchically, with volumes as the highest in the order and files as the lowest in the order. You can use the appropriate "Work with..." command to access these panels directly without having to go through the Optical Support Utilities main menu. Each display presents the selected information and the options that are available. Some options might not apply to all optical devices or volumes.

This chapter contains information describing which optical commands are supported to which devices and media formats. The different optical devices are:

- CD-ROM
- DVD-ROM
- DVD-RAM
- Directly attached optical media library
- LAN-attached optical media library
- Virtual optical device

The different optical media formats supported are:

- ISO 9660
- Universal Disk Format (subset of ISO 13346)
- High performance optical file system (HPOFS)

See Chapter 10, "Optical Media Formats," on page 113 for more information

Command and Device Dependency

Some optical commands have no meaning when used with certain optical devices. No support exists for other commands with certain optical devices. Table 4 lists all of the device related optical commands and the device types to which they apply.

Table 4. Optical Commands and Device Dependencies

Command	CD-ROM or DVD Device	Directly Attached Library Device	LAN-Attached Library Device	Virtual Device
ADDOPTCTG	Supported	Supported	Not supported	Not applicable
ADDOPTSVR	Not applicable	Not applicable	Supported	Not applicable
CHGDEVMLB	Not applicable	Supported	Not applicable	Not applicable
CHGDEVOPT	Supported	Not applicable	Not applicable	Supported
CHGOPTA ¹	Supported	Supported	Not supported	Supported
CHKOPTVOL	Supported	Supported	Not supported	Supported
CRTDEVMLB	Not applicable	Supported	Not applicable	Not applicable
CRTDEVOPT	Supported	Not applicable	Not applicable	Supported
DSPOPTSVR	Not applicable	Not applicable	Supported	Not applicable
RCLOPT	Supported	Supported	Not supported	Supported
RMVOPTCTG	Supported	Supported	Not supported	Supported

Table 4. Optical Commands and Device Dependencies (continued)

Command	CD-ROM or DVD Device	Directly Attached Library Device	LAN-Attached Library Device	Virtual Device
RMVOPTSVR	Not applicable	Not applicable	Supported	Not applicable
VFYOPT	Supported	Supported	Not supported	Not applicable
Notes:				
1. CHGOPTA is not a device or volume specific command.				

Working with Optical Volumes Option

The primary menu for working with optical volumes is the Work with Optical Volumes display (shown in Figure 14). There are several variations of this display to accommodate alternate formats and extended attribute information.

```

Work with Optical Volumes
System: XXXXXXXX
Device . . . . . *ALL
Side information . . . . . : *ALL
Type options, press Enter.
  1=Add  2=Change  3=Copy  4=Remove  5=Display  8=Work with directories
 10=Initialize 11=Work with object links 12=Duplicate ...

Opt  Volume      Device      Volume      Media      Authorization
   -  -----      -          Type        Type        List
   -  VOL2047      OPTMLB01   *PRIMARY   *ERASE      QOPTSEC
   -  VOL2048      OPTMLB01   *PRIMARY   *ERASE      QOPTSEC

Bottom

Parameters or command
====>
F3=Exit  F4=Prompt  F5=Refresh  F6=Print list  F9=Retrieve  F11=View 2
F12=Cancel  F14=Show extended information  F24=More keys ...

```

Figure 14. Work with Optical Volumes Display

You can select the Work with Optical Volumes display by choosing option 1 (Work with optical volumes) on the Optical Support Utilities menu. You can also run the Work with Optical Volumes (WRKOPTVOL) command on the command line.

Note: The Work with Optical Volumes (WRKOPTVOL) command applies to the following volumes:

- Volumes in CD-ROM or DVD devices.
- Volumes in directly-attached optical media library devices.
- Volumes in LAN-attached optical media library devices.
- Volumes in virtual optical devices

Displaying Optical Volumes

When the Work with Optical Volumes display first appears, it includes a list of all volumes in all CD-ROM devices, DVD devices, optical media libraries, and

LAN-attached devices. The volume names that are displayed are determined by what you type in the device (DEV) and CSI parameters. The following options are valid for the DEV parameter:

Option Explanation

- name* The name of a specific device. This lists all volumes in the specified device.
- *ALL** The list of all volumes in all devices. The volumes display in alphabetical order regardless of the device they are in.

You can press F11 (View 2) on the Work with Optical Volumes display to view the text variation of this display (shown in Figure 15).

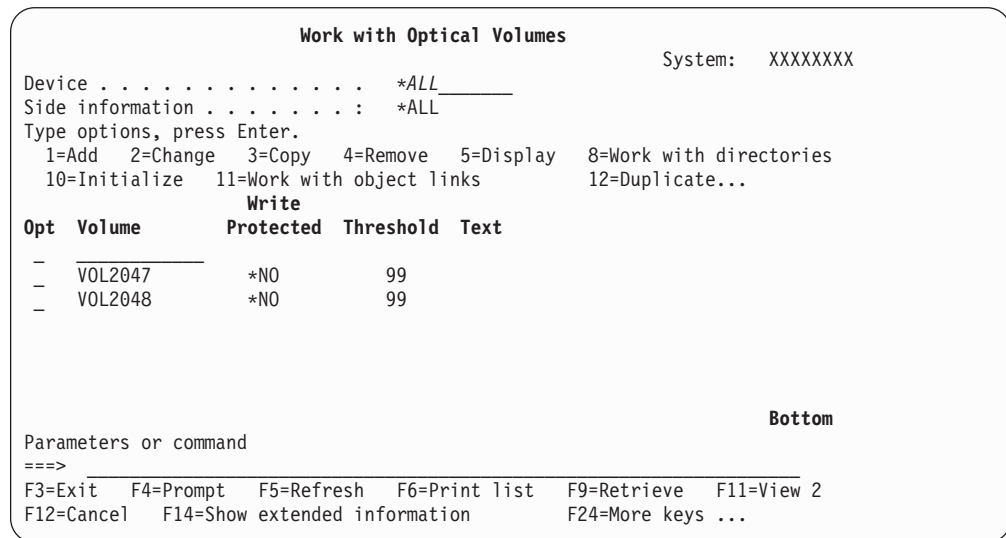


Figure 15. Work with Optical Volumes Display, View 2

Press F11 (View 1) to return to the status variation.

A third variation of the Work with Optical Volumes display is the extended information display. To view this display, press F14 (Show extended information) on the Work with Optical Volumes display, or use the WRKOPTVOL command and set the extended information parameter to *YES. This display is shown in Figure 16 on page 40.

```

Work with Optical Volumes
System: XXXXXXXX
Device . . . . . *ALL_____
Side information . . . . . : *ALL
Type options, press Enter.
1=Add 2=Change 3=Copy 4=Remove 5=Display 8=Work with directories
10=Initialize 11=Work with object links 12=Duplicate...

```

Opt	Volume	Location	% Used	Last Referenced	Volume Accessible
-	VOL2047	SLOT004	38.4	01/17/95	*NO
-	VOL2048	SLOT004	54.8	06/17/94	*NO

```

Bottom
Parameters or command
====>
F3=Exit F4=Prompt F5=Refresh F6=Print list F9=Retrieve F12=Cancel
F14=Exclude extended information F16=Repeat position to F24=More keys

```

Figure 16. Work with Optical Volumes Display, Extended Information

Unlike the status and text variations, this display requires that the optical device be varied on. If an optical device is not varied on, message OPT1520, Data displayed may not be current, is returned.

Initializing Optical Volumes

You must initialize the writable optical media before the system can create directories and files. When you initialize a volume, a new volume identifier must be given which gets written to the media.

You initialize an optical volume using the Initialize Optical command. To select this command from the Work with Optical Volumes display, select option 10 (Initialize) in the Opt (Option) column next to the volume you want to initialize. The Initialize Optical Volume display appears and prompts you for required information. Figure 17 on page 41 shows the Initialize Optical Volume display.

The Media Format parameter determines the media format of the volume. ***MEDTYPE** is the default which means that the media type determines the media format. The two media formats available are UDF (Universal Disk Format) and HPOFS (High Performance Optical File System). See Chapter 10, "Optical Media Formats," on page 113 for more details about these media formats and what formats are allowed on which media types.

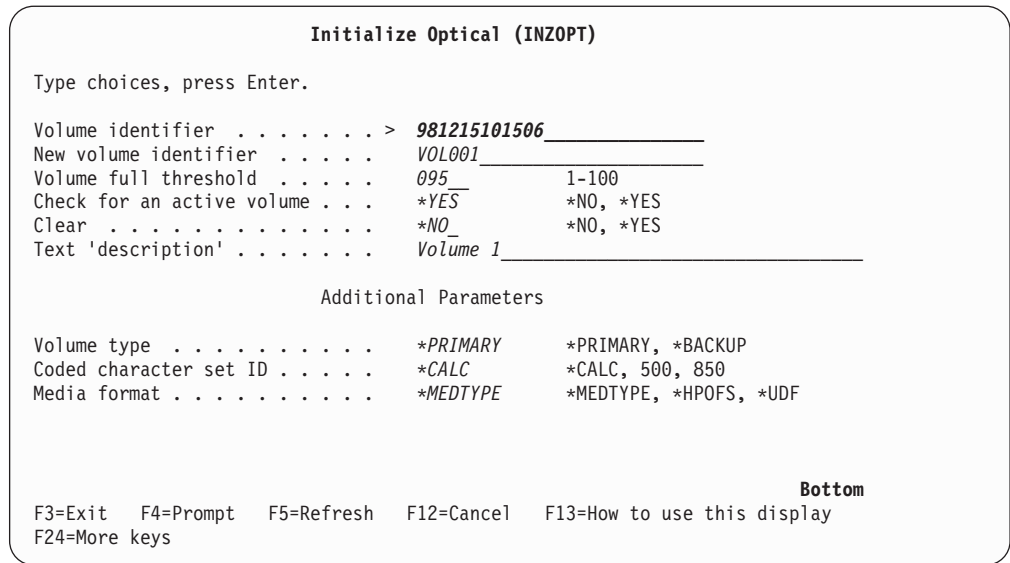


Figure 17. Initialize Optical (INZOPT) Display

Attention: When you initialize an optical volume, all information previously written on the volume becomes inaccessible.

Note: The Initialize Optical (INZOPT) Command applies to the following volumes:

- Volumes in directly-attached optical media libraries.
- Volumes on writable media in DVD devices.
- Volumes in virtual optical devices.

Renaming Optical Volumes

You can rename an optical volume without losing the information on the volume. The system merely updates the optical volume with the new name. To rename a volume, select option 7 (Rename) in the Opt (Option) column on the Work with Optical Volumes display.

The fields on this display show the following information for renaming optical volumes:

- Volume: The current name of the optical volume for renaming appears in this field.
- New Name: Specify the new name of the optical volume. The new name defaults to the current volume name.

Note: Beginning in Version 5 Release 2, you can rename unformatted (uninitialized and unknown) volumes. When an unformatted volume is renamed, the new name acts as an alias to the unformatted volume. The new name will not be written to the media and will not be preserved if the volume is exported from the device. The name is only a temporary volume identifier used to reference this volume.

Adding Optical Disk Cartridges

To add an optical disk cartridge, perform the following steps:

1. On the Work with Optical Volumes display, select option 1 (Add) in the Opt (Option) column next to the blank volume name. You are prompted by the Add Optical Cartridge display, shown in Figure 18 on page 42.




Figure 18. Add Optical Cartridge (ADDOPTCTG) Display

2. Enter information for the following:

- Optical media library (required).
- Authorization list.
- Rebuild directory index.

Note: Specifying *NO for Rebuild Directory Index can improve the performance of ADDOPTCTG by deferring the rebuilding of the optical directory index until a later time.

- Media location

Note: Some 399F models support adding multiple media through the bulk magazine. *MAGAZINE can be specified to use this feature if available for the device.

You do not need to provide a volume identifier. The system supplies a date and time stamp as the volume identifier. The date and time stamp is used to track each volume until it is read. If the volume is not initialized, the date and time stamp serves as its identifier until the volume is initialized with a user-supplied name.

Note: The Add Cartridge (ADDOPTCTG) command applies to:

- Directly-attached optical media libraries
- CD and DVD devices

Copying Optical Volume Data

To copy optical files in one or all of the directories on a volume to another volume or directory, select option 3 (Copy) in the Opt (Option) column on the Work with Optical Volumes display next to the volume with the directory you want copied. The Copy Optical (CPYOPT) display (shown in Figure 19 on page 43) appears and prompts you for more information.

```

                                Copy Optical (CPYOPT)

Type choices, press Enter.
From volume identifier:
  Volume . . . . . VOL001
  Volume type . . . . . *PRIMARY *PRIMARY, *BACKUP
  From path . . . . . /
  _____
  _____

To volume identifier:
  Volume . . . . . VOL004
  Volume type . . . . . *PRIMARY *PRIMARY, *BACKUP
  To path . . . . . *FROMPATH
  _____
  _____

Select files to copy . . . . . *CHANGED *CHANGED, *NEW, *ALL
Copy subdirectories . . . . . *NO *NO, *YES
                                                    More...

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

```

```

                                Copy Optical (CPYOPT)

Type choices, press Enter.

Additional Parameters

Create directory . . . . . *NO *NO, *YES
Allow copy to opposite side . . *NO *NO, *YES
Copy option . . . . . *IOP__ *IOP, *SYSTEM
Starting date and time:
  Starting date . . . . . *BEGIN__ Date, *BEGIN
  Starting time . . . . . *AVAIL__ Time, *AVAIL

                                                    Bottom

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys
Parameter FROMVOL required.

```

Figure 19. Copy Optical (CPYOPT) Display

This command is not allowed for LAN-attached optical devices. If you attempt to use a LAN-attached optical device, an error message is issued.

The name of the volume appears on the display. You must specify the name of the directory to be copied from and the volume to receive the copy (the volume must be initialized before copying). If you are copying a full volume, specify an unused volume to receive the copied files to be sure that enough space is available.

This command does not delete files for you. Therefore, if you use this command to do additional copies, you must delete files from the target volume that have been deleted from the source volume.

When the copy request completes, a message is added to the job log stating the number of files copied successfully and the number of files that were not copied. For each file that is not copied, a message stating the full file name is added to the

job log. For each directory processed, a directory copied message is added to the job log stating the number of files copied successfully and the number of files that were not copied successfully.

Select Files to Copy (SLTFILE) Parameter

The Select files to copy (SLTFILE) parameter indicates how files are to be selected for copying. You can select whether to replace files that already exist on the volume to which you are copying. A value of *CHANGED specifies that a file is copied if does not exist on on the target volume, or if the file is more current than the one on the target volume. A value of *NEW specifies that only files that do not already exist on the To volume identifier field are copied. A value of *ALL specifies that all files are copied, even if they exist with the same creation date.

Copy Option (COPYTYPE) Parameter

The Copy option (COPYTYPE) parameter indicates which resources are used to perform the copy operation. A value of *IOP specifies that the copy operation will have better performance but will slow down other requests to the optical media library. A value of *SYSTEM specifies that the copy request will share the optical media library resources with other requests but will cause the copy request to take longer.

Starting Date and Time (FROMTIME) Parameter

The starting date and time (FROMTIME) parameter is optional and can be used to further restrict the number of files that are copied. Files from the source volume that have a creation or modification date and time greater than or equal to the date and time entered on this parameter are selected to be copied.

Copying Optical Volume Data—Example

To make a complete copy of VOLA on VOLB, use the following command:

```
CPYOPTDIR FROMVOL(VOLA) FROMDIR('/') TOVOL(VOLB) TODIR(*FROMDIR)
SLTFILE(*CHANGED) CPYSUBDIR(*YES) CRTDIR(*YES) ALWCPYOPP(*NO)
COPYTYPE(*IOP)
```

Note: The Copy Optical (CPYOPT) command applies to:

- Volumes in directly-attached optical media libraries
- Volumes in CD-ROM or DVD devices
- Volumes in virtual optical devices

Changing Optical Volume Attributes

You can change the attributes of a volume by typing a 2 (Change) in the Opt (Option) column of the Work with Optical Volumes display. The Change Optical Volume (CHGOPTVOL) display (shown in Figure 20 on page 45) appears and prompts you to type your changes for the volume you selected.

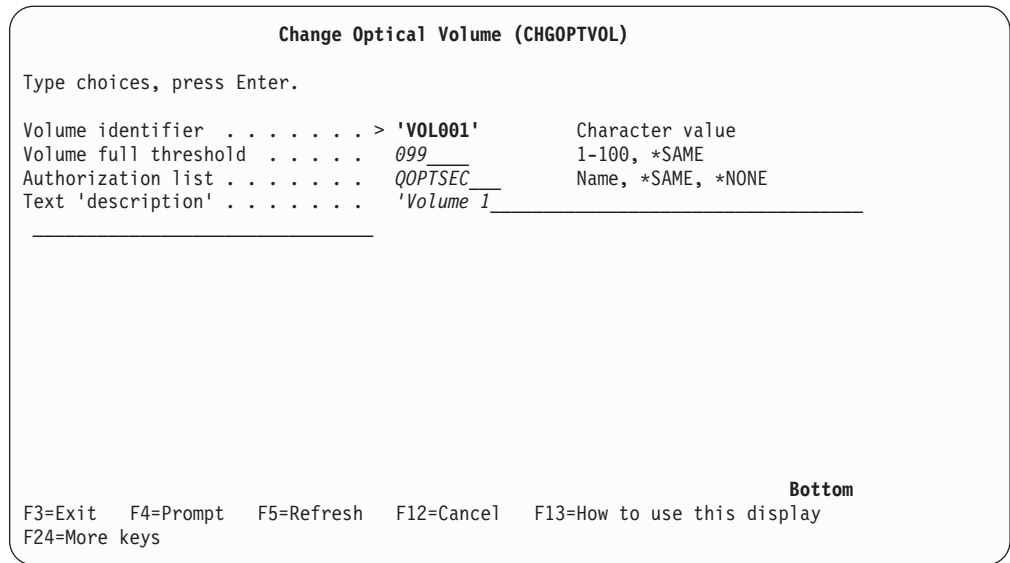


Figure 20. Change Optical Volume (CHGOPTVOL) Display

You can change the following attributes with this command:

- Volume-full threshold percentage for volumes in directly-attached media libraries.
- Authorization list that is used to secure the volume.

Note: If the volume is in a stand-alone device (CD-ROM or DVD drive), the authorization list secures the volume for the duration that the media is in the device. Ejecting the media from the device and immediately re-inserting it again resets the authorization list to the QOPTSEC default. The system does not maintain the authorization list for volumes that are removed from a stand-alone optical device! For volumes in an optical media library device, you can only maintain the authorization list when removing the media by specifying VOLOPT(*KEEP) on the Remove Optical Cartridge (RMVOPTCTG) CL command. The authorization list is not written to the optical disk but instead is maintained internally on the server.

- Volume description of the volume for DVD-RAM volumes and volumes in directly-attached media libraries.

Displaying and Printing Optical Volume Attributes

Displaying Optical Volume Attributes

To view the attributes of a volume, select option 5 (Display) in the Opt (Option) column next to the volume you want to view on the Work with Optical Volumes display. The first Display Optical Volume Attributes display appears (Figure 21 on page 46). This display shows the attributes of the volume you selected. All fields may not contain a value depending on the type of volume and where the volume is located.

```

Display Optical Volume Attributes

Volume . . . . . : VOL001
Device . . . . . : OPT02
Authorization list . . . . . : QOPTSEC

Volume Header Information:
Internal Volume ID . . . . . : VOL001
Serial Number . . . . . : 1077952576
Volume type . . . . . : *PRIMARY
Media type . . . . . : *ERASE
Media format . . . . . : *HPOFS
Coded character set ID . . . . . : 500
Volume-full threshold . . . . . : 99
Volume sequence number . . . . . : 0
Create date . . . . . : 12/28/94
Create time . . . . . : 15:43:46
Text . . . . . : Annual sales volume by state

More...

Press Enter to continue.

F3=Exit F12=Cancel F14=Additional volume attributes

```

Figure 21. Display Optical Volume Attributes Display, First Display

Press the Page Down key to view the second display (Figure 22). The second Display Optical Volume Attributes display appears.

```

Display Optical Volume Attributes

Usage Information:
Last reference date . . . . . : 01/24/95
Volume on opposite side . . . . . : VOL002
Block size . . . . . : 1024
Volume capacity (bytes) . . . . . : 305135616
Space available on volume (bytes) . . : 188000000
Percentage used . . . . . : 38.38

Status Information:
Volume location . . . . . :
Volume access . . . . . : Writeable
Double volume . . . . . : No
Double sided . . . . . : No
IPL Capable . . . . . : No
Volume held by . . . . . : *NONE

Bottom

Press Enter to continue.

F3=Exit F12=Cancel F14=Additional volume attributes

```

Figure 22. Display Optical Volume Attributes Display, Second Display

You cannot change any information on either of these displays.

A third display is possible if the optical volume type is ***BACKUP**. If the display indicates More... in the bottom right-hand corner, press the Page Down key to view the third display. (This display shows information that is unique to optical backup volumes.)

Press F14 to see the Additional Volume Attribute displays.

Printing Optical Volume Attributes

You can print volume attributes by typing a 6 (Print) in the Opt (option) column next to a volume listed on the Work with Optical Volumes display.

The output is written to a spooled file, which can be found in the output queue for the job.

Note: The Display Optical (DSPOPT) command applies to:

- Volumes in optical CD-ROM or DVD media devices.
- Volumes in directly-attached optical media libraries.
- Volumes in optical LAN-attached devices.
- Volumes in virtual optical devices

Duplicating Optical Volumes

An efficient method to create a backup of an optical volume is to use the Duplicate Optical (DUPOPT) CL command. This command performs sector copies to create a volume which is identical to the source except for the volume identifier and creation date and time.

Refer to Chapter 8, “Optical Volume Backup,” on page 89 for more information about using the Duplicate Optical command.

Duplicate Optical (DUPOPT)

Type choices, press Enter.

From volume identifier	<u>VOL001</u>	_____
To volume identifier	<u>980428072748</u>	_____
New volume identifier	<u>VOL001-BACKUP</u>	_____
Clear	<u>*NO_</u>	<u>*NO, *YES</u>
From device	_____	Name,
To device	_____	Name,

Additional Parameters

From end of media option	<u>*LEAVE_</u>	<u>*LEAVE, *UNLOAD</u>
To end of media option	<u>*LEAVE_</u>	<u>*LEAVE, *UNLOAD</u>

Bottom

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

Figure 23. DUPOPT Panel

Enter the information for the following fields:

- From volume identifier.
- To volume identifier.
- New volume indentifier.
- Clear.

The Duplicate Optical (DUPOPT) command applies to the volumes in directly-attached optical media library devices and DVD devices.

Working with Directories and Files

There are two methods to view directory and file information through the optical support panels and commands.

Work with Object Links

Work with Object Links (WRKLNK) works with directories and files. This command gives a PC-like hierarchical view of the directories and files on the volume. Both directories and files at the given level in the path hierarchy will display. The system denotes directories as type DDIR and files as type DSTMF.

```
Work with Object Links

Directory . . . . : /QOPT/VOL001

Type options, press Enter.
  2=Edit  3=Copy  4=Remove  5=Display  7=Rename  8=Display attributes
  11=Change current directory ...

Opt  Object link      Type  Attribute  Text
--  -
--  DIRECTORY.1      DDIR
--  DIRECTORY.2      DDIR
--  DIRECTORY.3      DDIR
--  FILE.1           DSTMF
--  FILE.2           DSTMF
--  FILE.3           DSTMF

Parameters or command
===>
F3=Exit  F4=Prompt  F5=Refresh  F9=Retrieve  F12=Cancel  F17=Position to
F22=Display entire field  F23=More options

Bottom
```

Figure 24. WRKLNK Panel

The Work with Object Links (WRKLNK) command applies to the following conditions:

- Volumes in CD-ROM or DVD devices.
- Volumes in directly-attached optical media library devices.
- Volumes in virtual optical devices

Work with Optical Directories and Files

Work with Optical Directories (WRKOPTDIR) command works only with directories. You can display all directories and subdirectories, or just display certain levels if desired. This command requires creating the optical directory index if it was not created during Add Optical Cartridge. The Work with Optical Files (WRKOPTF) command works with optical files.

The Work with Optical Directories (WRKOPTDIR) command and Work with Optical Files (WRKOPTF) command apply to the following conditions:

- Volumes in directly-attached optical media library devices.
- Volumes in LAN-attached optical media library devices.
- CD-ROM volumes in either CD-ROM or DVD devices.

Note: Volumes that are created in Universal Disk Format do not support the WRKOPTDIR and WRKOPTF commands.


```

                                Work with Optical Directories
                                System:  XXXXX
Directory . . . . . *ALL
Volume . . . . . VOL001
Date . . . . . *AVAIL__

Type options, press Enter.
  1=Create  3=Copy  4=Delete  5=Display  6=Print  8=Work with files

Opt Directory                                Volume      Created
- /DIRECTORY.1/SUB1/SUB2                       VOL001      04/05/99
- /DIRECTORY.1/SUB1                             VOL001      04/05/99
- /DIRECTORY.2                                  VOL001      04/05/99
- /DIRECTORY.1                                  VOL001      04/05/99
- /                                              VOL001      04/05/99

Bottom

Parameters or command
===>
F3=Exit  F4=Prompt  F5=Refresh  F6=Print list  F9=Retrieve
F11=Sort by directory  F12=Cancel  F16=Repeat position to  F24=More keys

```

Figure 25. WRKOPTDIR Panel

Removing Optical Volumes

To remove a volume, select option 4 (Remove) in the Opt (Option) column next to the volume you want to remove on the Work with Optical Volumes display. You can then remove an optical volume by physically removing the optical disk cartridge from the optical library dataserwer. The Remove Optical Cartridge display (Figure 26) appears and prompts you for additional information.

```

                                Remove Optical Cartridge (RMVOPTCTG)

Type choices, press Enter.

Volume identifier . . . . . > 'VOL001'          Character value, *IOSTATION
Volume description option . . . *REMOVE         *REMOVE, *KEEP
Media location . . . . . *IOSTATION            *IOSTATION, *MAGAZINE

Bottom

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys

```

Figure 26. Remove Optical Cartridge (RMVOPTCTG) Display

Enter the information for the following fields:

- Volume identifier
- Volume description option

- Removed cartridge location (if *KEEP was specified as the volume description option)
- Media location: Some 399F models support removing media to the bulk magazine. *MAGAZINE can be specified to use this feature if available for the device. The magazine can be removed by specifying VOL(*MAGAZINE) on this command.

Because there are two volumes on each optical disk, the options selected on the Remove Optical Cartridge display apply to both volumes.

Note: The Remove Cartridge (RMVOPTCTG) command applies to:

- Volumes in directly-attached optical media libraries
- Volumes in CD-ROM and DVD devices
- Volumes in virtual optical devices

Deleting *REMOVED Volumes from the Optical Index Database

If you removed the volume but saved the volume description information, you can later delete that information by selecting option 9 (Delete). The delete option can also be used if a volume has been marked as being in a offline device. The delete option removes a single volume, and not both volumes of an optical cartridge.

Note: The delete option applies to:

- Removed volumes from an optical media libraries.
- Volumes in an offline optical media libraries.
- Volumes in an offline optical LAN device.

Check Optical Volume

To validate the integrity of the directories and files on the volume, select option 13 (Check) in the Opt (Option) column next to the volume identifier. This command can be used to verify all files can be read and will print a list of damaged files if any exist. In addition, a count of damaged and not damaged files can be output depending on the value of the OUTPUT parameter.

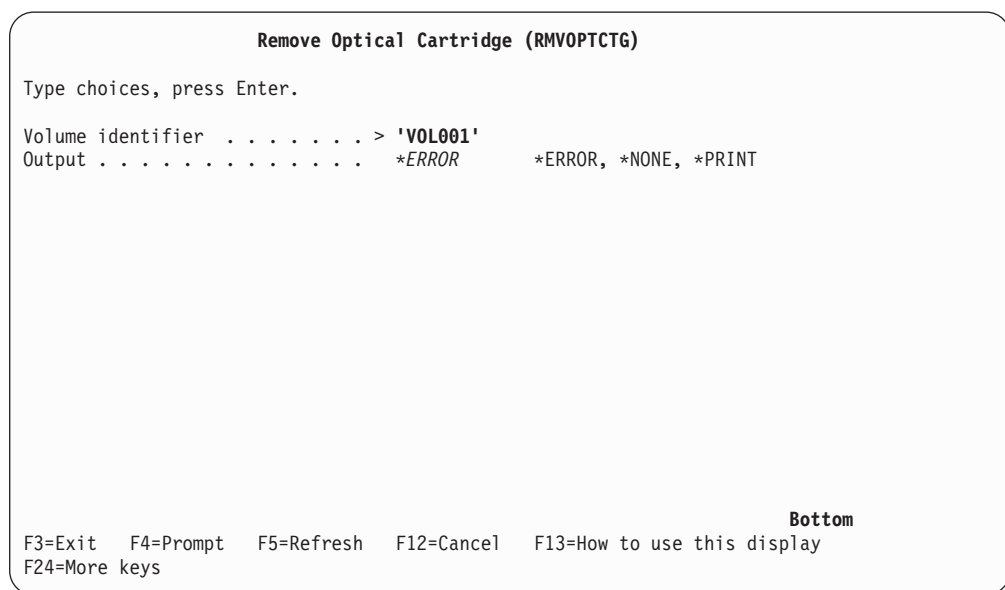


Figure 27. Check Optical Volume (CHKOPTVOL) Display

| Enter the information for the following fields:

- | • Volume identifier
- | • Output: Specifies conditions to print output
- | • Optical device: Required only if Volume identifier is *MOUNTED

| **Note:** The Check Optical Volume (CHKOPTVOL) command applies to:

- | • Volumes in directory-attached optical media libraries
- | • Volumes in CD-ROM and DVD devices
- | • Volumes in virtual optical devices

Chapter 6. Optical Security and Auditing

You can secure information on optical media by using OS/400 security functions. The level of security available depends on the optical media format of the volume. You can use an authorization list to secure all optical volumes. This includes all volumes in CD-ROM, DVD, LAN-attached, directly-attached, and virtual optical devices. Optical volumes formatted in Universal Disk Format (UDF) provide directory and file level security in addition to authorization list security.

Optical support provides ways to prevent unauthorized access and processing of data that is stored on optical volumes. Optical support does this by verifying a requester's rights to specific optical volumes before attempting the following requests:

- Open file or directory
- Create directory
- Delete file or directory
- Rename file
- Initialize or rename volume
- Remove cartridge
- Change or retrieve attributes
- Copy
- Backup or convert backup
- Save or release held files
- Read sectors
- Save optical volume storage
- Restore optical volume storage
- Check optical volume for damaged files

Along with security for optical volumes, directories, and files; auditing of access to optical objects is also available.

Authorization List Security

You can use an authorization list to secure all optical volumes. The authorization list contains the definition for each user's access to the volume. The system assigns an authorization list to the volume when the volume is added to the system. You can change the authorization list to the volume by using the Change Optical Volume (CHGOPTVOL) CL command. IBM ships the default authorization list that is named QOPTSEC with the iSeries server. This authorization list secures volumes unless you specifically secure the volume by another authorization list. You can change this, but if you eject and add the media, the system resets the authorization list to QOPTSEC. To add volumes to directly-attached media library devices, specify an authorization list on the Add Optical Cartridge command.

You can change the list of users in the default authorization list or the public authority to a volume. You can do this by using the authorization list management commands of OS/400. You must have authorization list management (*AUTLMGT) authority to change the users in an authorization list. For a complete description of authorization list management rules see the iSeries Security Referencebook.

Note: The authorization list is not written to the optical disk, but instead is maintained internally on the server. Therefore, the authorization list for volumes in stand-alone devices is not preserved when the media is removed and added to the device. For library devices, the authorization list can be preserved by using VOLOPT(*KEEP) on the Remove Optical Cartridge (RMVOPTCTG) CL command.

Authorities Required for Optical Functions

You must have *USE authority to an optical volume to use the following optical functions:

- Open file for read
- Open directory
- Retrieve file or directory attributes
- Read sector
- List paths or files
- Check optical volume for damaged files

You must have *CHANGE authority to an optical volume to use the following optical functions:

- Open for write or read write
- Create or delete directory
- Create, delete, or rename file
- Change file or directory attributes
- Save or release held optical file

You must have *ALL authority to an optical volume to use the following optical functions:

- Initialize volume (requires *CHANGE for DVD-RAM media).
- Rename volume (requires *CHANGE for DVD-RAM media).
- Convert backup volume to primary
- Duplicate optical volume (requires *CHANGE for DVD-RAM media).

You must have *USE authority to the source optical volume and *CHANGE authority to the target optical volume to use the following optical functions:

- Copy file
- Copy directory

You must have *CHANGE authority to the source optical volume and *CHANGE authority to the target optical volume to use the following optical functions:

- Move file

You must have *AUTLMGT authority to an optical volume to use the following optical functions:

- Change authorization list used to secure the volume
- Add optical cartridge (If overriding existing authorization list)

You must have *OBJEXIST authority to the source volume to use the following optical functions.

- Save optical volume

You must have *OBJEXIST authority to the target volume to use the following optical functions.

- Restore optical volume

Specifying an Authorization List on the Add Optical Cartridge (ADDOPTCTG) Command

The authorization list (AUTL) parameter on the Add Optical Cartridge (ADDOPTCTG) command allows the volumes being imported into an optical media library to be automatically secured with an authorization list as part of the import processing. If no authorization list is specified for a new optical volume (one that has not been removed with the VOLOPT(*KEEP) option), the default optical authorization list (QOPTSEC) is used to secure the volume. If the volume was removed with the VOLOPT(*KEEP) option, the authorization list that previously secured the volume is used to secure the volume.

The user that adds the optical cartridge does not need to have any authority to the data on the volume being secured by the authorization list, as long as the user is not overriding the authorization list that previously secured a volume removed with the VOLOPT(*KEEP) option.

This is different from the way authorization lists are used to secure system objects. For example, a system operator should be able to add the PAYROLL optical disk to the optical media library and secure it with the PAYROLL authorization list, but not be able to access the data on the PAYROLL optical disk.

To change the authorization list used to secure an optical volume previously removed with the VOLOPT(*KEEP) option, the user issuing the ADDOPTCTG command must have either *AUTLMGT authority to the authorization list that previously secured the volume or *ALLOBJ special authority.

Authorization List Assignment and the Add Optical Server (ADDOPTSRV) Command

The Add Optical Server (ADDOPTSRV) command secures all volumes in the server with the default optical authorization list (QOPTSEC) unless a volume was previously secured with a different authorization list and then removed using the VOLOPT(*KEEP) option on the Remove Optical Server (RMVOPTSVR) command.

Changing the Authorization List to Secure an Optical Volume

You can change the authorization list that is used to secure an optical volume in a CD-ROM device or a directly or LAN-attached optical media library by using the Change Optical Volume (CHGOPTVOL) command. The user attempting to change the authorization list that is used to secure an optical volume must have either *AUTLMGT authority in the authorization list that is currently securing the volume or *ALLOBJ special authority. Specifying *NONE as the authorization list name causes access verification processing for the volume to be bypassed on the future access requests for the volume. If the new authorization list does not exist, the CHGOPTVOL command is rejected and a message is issued indicating that the new authorization list does not exist. Whenever the authorization list used to secure a volume is changed to a different authorization list or to *NONE, an audit entry is logged if optical auditing is active.

Retaining the Authorization List when Removing Optical Volumes

The tie between an optical volume and an authorization list is maintained in the optical index database. This tie is lost when a volume is exported with the *REMOVE option, because the record is deleted.

If the *KEEP option is specified when the volume is removed, the record is kept. By specifying *PREV on the Add Optical Cartridge (ADDOPTCTG) command, the authorization list which secured the optical volume before it was removed with the *KEEP option is used to secure the volume when it is re-added. The tie between a CD-ROM volume and the authorization list securing it is lost when the CD-ROM is removed from the drive.

Mapping an Authorization List to an Optical Volume

The name of the authorization list used to secure an optical volume is kept in an optical index database file. If an authorization list used to secure an optical volume cannot be found when attempting to access the volume, the access is denied and a message is issued indicating that the authorization list for the volume could not be found. If *NONE is specified as the authorization list used to secure an optical volume, no access verification is performed. The authorization list that secures an optical volume can be determined by using the Work with Optical Volumes (WRKOPTVOL) command.

Directory and File Level Security

Directory and file level security is available for volumes formatted in Universal Disk Format (UDF). The system maintains the "data authorities" of optical directories and files for three groups of users — owner, group, and public. You can display, change, and manage these authorities by using the integrated file system authority commands Display Authority (DSPAUT), Change Authority (CHGAUT), and Work with Authority (WRKAUT). To change the owner and primary group for files and directories, use the integrated file system commands Change Owner (CHGOWN) and Change Primary Group (CHGPGP). You can also access these commands through the Work with Optical Volumes display by selecting **option 11 (Work with object links)** on the desired volume.

For detailed information on the directory and file authorities that are required for specific interfaces, refer to Chapter 7, "Using Application Programming Interfaces for Optical Support," on page 59. Also refer to Chapter 10, "Optical Media Formats," on page 113 for an overview of the media format that provides directory and file level security.

Optical Auditing

To enable optical auditing, the system value QAUDCTL must be set to *AUDLVL and *OPTICAL must be specified in the QAUDLVL system value. You use the *SEC value on the SYSVAL parameter of the Work with System Values (WRKSYSVAL) command to change these system values. Changing these values requires *AUDIT special authority.

The following optical operations can be audited:

- Create, copy, or delete a directory
- Open file, including access mode (read only, write only, read and write)
- Copy, move, rename, or delete file

Change or retrieve directory attributes
Control file system (save or release held file, sector read) options
Open a directory
Back up optical volumes
Initialize or rename an optical volume
Convert a backup optical volume to a primary volume
Add or remove an optical cartridge
Change the authorization list, securing an optical volume
Saving an optical volume
Restoring an optical volume
Check optical volume for damaged files

The format of the optical audit journal entries can be found in the iSeries Security Reference book.

Object Authority

All programs are shipped with PUBLIC(*EXCLUDE) authority, and most of the commands are shipped with PUBLIC(*USE) authority. The following commands are shipped PUBLIC(*EXCLUDE).

- Add Optical Cartridge (ADDOPTCTG)
- Remove Optical Cartridge (RMVOPTCTG)
- Add Optical Server (ADDOPTSVR)
- Remove Optical Server (RMVOPTSVR)
- Reclaim Optical (RCLOPT)

Refer to the iSeries Security Reference book for information on controlling access to commands.

Chapter 7. Using Application Programming Interfaces for Optical Support

This chapter describes the interfaces an application programmer may use to access optical volumes on OS/400. The proper name for these interfaces is application programming interfaces, or APIs. You can use these APIs to interface with a variety of different file systems, of which the optical file system is one. This chapter is not a complete description of these APIs. It only describes the aspects of API usage that are unique to the optical file system.

Programmers can utilize the available APIs to create, access, change, or maintain optical files and directories. The APIs may be used to customize usage of optical support for specific business applications.

Categories of APIs for Optical Support

Two categories of APIs may be used to manipulate optical files and directories:

- Hierarchical file system (HFS) APIs
- Integrated file system support, which consists of UNIX-type APIs and the generic command interface

You can use both categories of APIs concurrently. For example, an optical file that is opened for reading by one application by using the HFS Open Stream File API can be opened for reading by another application by using the Open UNIX-type API. To find out more information about using the HFS or UNIX-type APIs on the generic commands, refer to the APIs topic in the Programming category of the Information Center.

Since different file systems exist in OS/400, some means must be provided for HFS or the integrated file system to differentiate which file system a call is targeted for. This is accomplished by requiring that the first name in the path name parameter be the name of the file system to be called, preceded by a leading slash. For the optical file system to be identified as the receiver of a request submitted to HFS or the integrated file system, the first portion of the path name parameter must be /QOPT.

The following topics describe the aspects of HFS or integrated file system API usage that are unique to the optical file system.

Using the HFS API Interface to Optical Support

A hierarchical file system (HFS) is a part of the operating system that includes the application programming interface (API) and the underlying file system (optical or otherwise) support. The HFS API makes it possible for an application that is written in a high-level language to create, store, retrieve, and manipulate data on a Directly-attached optical library device, LAN-attached optical library device, CD-ROM, or DVD device. To find out more information about HFS API refer to the refer to the APIs topic in the Programming category of the Information Center.

The HFS API support for optical support consists of two parts:

- An application programming call interface to the hierarchical file system to manipulate objects known as files and directories.

- An optical or other registered file system that manages the storage devices where the files and directories are stored.

HFS API optical functions include the following:

- Creating or deleting a directory
- Opening, reading, or closing a directory
- Opening, reading, writing, or closing a file
- Locking or unlocking bytes in a file
- Getting or setting the size of a file
- Renaming, copying, deleting, or removing a file
- Retrieving or changing directory entry attributes

Applications use HFS APIs to manage stream files on an iSeries server. These stream files are also called objects to identify them as data elements that do not have a conventional record structure. The object is treated as a named byte stream of known length, whose size can vary from a few bytes to megabytes.

HFS APIs allow applications to create and manage file objects on storage devices and to perform input/output operations to those file objects. HFS APIs allow applications to create and manage directory objects, which can be thought of as a logical grouping of similar file objects. These directory objects contain information about the file objects belonging to that directory. Directories can be contained within directories resulting in the hierarchical structure.

Optical File System Implementation of HFS APIs

This topic describes how the HFS API usage is different for the optical file system as compared to the specifications in the **Information Center**. Although the APIs that HFS supports are common to all file systems, each file system has different interpretations or restrictions regarding those APIs. Table 5 summarizes the optical interpretation of each HFS API. LAN-attached optical devices and directly-attached optical devices have different restrictions for several of the APIs.

Table 5. Optical HFS API Restrictions

HFS APIs	Directly Attached Usage Notes	LAN-Attached Usage Notes
Change File Pointer (QHFCHGFP)	None	None
Close Stream File (QHFCLOSF)	None	None

Table 5. Optical HFS API Restrictions (continued)

HFS APIs	Directly Attached Usage Notes	LAN-Attached Usage Notes
Control File System (QHFACTLS)	<p>Supports the following requests:</p> <ul style="list-style-type: none"> • SAV saves a held optical file. • RLS releases a held optical file. • SRD/VOL returns a sector read from an optical volume. • SRD/DEV returns a sector read from an optical device. • RTV/VOL returns volume specific information. • GET reads file data directly from the media with minimal data caching. For UDF formatted volumes, GET requires *X authority to each directory in the path preceding the file and *R authority to the file. 	<p>Supports the following requests:</p> <ul style="list-style-type: none"> • UPD/LAN performs a dynamic index refresh of the list of LAN volumes • UPD/VOL returns volume specific information • RTV/VOL returns volume specific information • RTV/DIR returns subdirectory and file entries for a specified directory
Copy Stream File (QHFCPYSF)	<p>If the source file is in the QOPT file system, *USE authority is required to the source optical volume.</p> <p>If the target file is in the QOPT file system, *CHANGE authority is required to the target optical volume. Copy information parameter, Byte 1, Option 2 is not supported (Copy Append). If specified, CPF1F62 will be returned.</p> <p>When the operation is complete, QCRTDTM, QACCDTM, and QWRDTM are set to the current date.</p> <p>When copying between the QOPT and QDLS file systems, file attributes are optionally copied depending on global optical attribute CPYATR. This attribute can be displayed or changed utilizing the CHGOPTA command.</p> <p>When copying between the QOPT and QDLS file systems, file permissions are not copied. If permissions need to be preserved between these file systems use the copy (CPY) CL command.</p> <p>If the source file is on a UDF formatted volume, *X authority is required to each directory in the path preceding the file. *R authority is required to the file.</p> <p>If the target file is on a UDF formatted volume, *WX authority is required to the parent directory and *X authority to each directory in the path preceding the parent directory.</p>	<p>If the source file is in the QOPT file system, *USE authority is required to the source optical volume.</p> <p>If the target file is in the QOPT file system, *CHANGE authority is required to the target optical volume. Copy information parameter, Byte 1, Option 2 is not supported (Copy Append).</p> <p>Copying from a volume in a directly attached library to a volume in a LAN-attached optical device is not supported.</p>

Table 5. Optical HFS API Restrictions (continued)

HFS APIs	Directly Attached Usage Notes	LAN-Attached Usage Notes
<p>Create Directory (QHFCRTDR)</p>	<p>When the operation is complete, QCRTDTTM, QACCDTTM, QWRDTTM are set to the current date.</p> <p>When the operation is complete, QFILSIZE and QALCSIZE are set to 0.</p> <p>Requires *CHANGE authority to the optical volume.</p> <p>Creating the optical root directory is not supported.</p> <p>Creating the volume portion of the directory is not supported.</p> <p>Attributes passed in the attribute information table are not supported, and will result in a CPF1F71 error message. The length of the attribute information table parameter must be 0.</p> <p>Optical attribute OPT.CHGATDTTM, which indicates the last time that the directory attributes were changed, is created. This date is set to the current date. If a user specifies an attribute, it is ignored.</p> <p>For UDF formatted volumes, *WX authority is required to the parent directory. *X authority is required to each directory in the path preceding the parent directory. The owner of the directory will be the user creating the directory and the owner data authorities will be set to *RWX. The primary group and primary group data authorities will be the same as the parent directory. The *PUBLIC data authorities will be the same as the parent directory.</p>	<p>When the operation is complete, QCRTDTTM, QACCDTTM, QWRDTTM are set to the current date.</p> <p>When the operation is complete, QFILSIZE and QALCSIZE are set to 0.</p> <p>Requires *CHANGE authority to the optical volume.</p> <p>Creating the optical root directory is not supported.</p> <p>Creating a volume portion of a directory is not supported.</p> <p>All standard attributes are ignored.</p> <p>Length of attribute information table parameter must be set to 0.</p>
<p>Delete Directory (QHFDLDR)</p>	<p>Deleting the optical root directory is not supported.</p> <p>Deleting the volume portion of a path is not supported.</p> <p>Requires *CHANGE authority to the optical volume.</p> <p>For UDF formatted volumes, *WX authority is required to the parent directory and *X authority is required to each directory in the path preceding the parent directory. *W authority is required to the directory being deleted.</p>	<p>Deleting the optical root directory is not supported.</p> <p>Deleting the volume portion of a path is not supported.</p> <p>Requires *CHANGE authority to the optical volume.</p>

Table 5. Optical HFS API Restrictions (continued)

HFS APIs	Directly Attached Usage Notes	LAN-Attached Usage Notes
Delete Stream File (QHFDLTSF)	Requires *CHANGE authority to the optical volume. For UDF formatted volumes, *WX authority is required to the parent directory. *X authority is required to each directory in the path preceding the parent directory. *W authority is required to the file being deleted.	Requires *CHANGE authority to the optical volume.
Get File Size (QHFGETSZ)	None	None
Set File Size (QHFSETSZ)	None	Not Supported

Table 5. Optical HFS API Restrictions (continued)

HFS APIs	Directly Attached Usage Notes	LAN-Attached Usage Notes
<p>Open Stream File (QHFOFNSF)</p>	<p>Parameter open information:</p> <ul style="list-style-type: none"> • Opening with an access mode (Byte 6) of write only or read/write requires *CHANGE authority to the volume. • Opening with an access mode (Byte 6) of read only requires *USE authority to the volume. • Lock Modes (Byte 5) are enforced across different open instances. If the same job opens a file multiple times, these open locks can conflict. <p>If QALCSIZE was specified on an open request for the write operation, optical media will be checked to see if enough space is available. If not, error message CPF1F62 is returned.</p> <p>All standard attributes except QALCSIZE are ignored.</p> <p>If a file is being created, QCRTDTM, QACCDTM, and QWRDTM are set to the current date. If a file is being updated, QWRDTM is set to the current date. If a file is being read, no time stamps are changed. QACCDTM is never changed after a file is created. It will always equal QCRTDTM.</p> <p>The following authorization rules apply only for UDF formatted volumes.</p> <ul style="list-style-type: none"> • If opening a file for READ, *X authority is required to each directory in the path preceding the file and *R authority is required to the file. • If opening an existing file for WRITE, *X authority is required to each directory in the path name preceding the file and *W authority is required to the file. • If opening an existing file for READ/WRITE, *X authority is required to each directory in the path name preceding the file and *RW authority is required to the file. • If creating the file, *WX authority is required to the parent directory. • If creating the file, the owner of the file will be the user creating the file and the owner data authorities will be set to *RWX. The primary group and primary group data authorities will be the same as the parent directory. The *PUBLIC data authorities will be the same as the parent directory. 	<p>Parameter Open information:</p> <ul style="list-style-type: none"> • Byte 3 (write-through flag), is not supported. • Byte 7 (type of open operation to perform), is not supported. • Opening with an access mode (Byte 6) of read only requires *USE authority to the volume. <p>Unless the file open attempt is for read only access, attributes are not tolerated and result in a CPF1F71 error message. The length of the attribute information table parameter must be 0.</p> <p>If a file open attempt is for read only access, attributes are tolerated but ignored.</p>

Table 5. Optical HFS API Restrictions (continued)

HFS APIs	Directly Attached Usage Notes	LAN-Attached Usage Notes
Read Stream File (QHFRDSF)	None	None
Retrieve Directory Entry Attributes (QHFRTVAT)	<p>Requires *USE authority to an optical volume.</p> <p>For UDF formatted volumes, *X authority is required to each directory in the path name preceding the file and *R authority is required to the file or directory being read.</p>	<p>User can retrieve only LAN-standard attributes: QFILSIZE, QCRTDTTM, and QWRDTTM.</p> <p>Requires *USE authority to an optical volume.</p> <p>Length of attribute information table parameter must be set to 0.</p>
Write Stream File (QHFWRTSF)	None	None
Change Directory Entry Attributes (QHFCHGAT)	<p>QFILATTR is the only standard attribute that can be changed. All others that are specified are ignored.</p> <p>Read only flag, byte 1 of the QFILATTR attribute, can only be set for a file, not a directory. If specified for a directory, it is ignored.</p> <p>Changed Flag, byte 5 of the QFILATTR attribute, can be set to either 0 or 1. It is automatically set on (1) whenever a file is created or written to.</p> <p>If OPT.CHGATDTTM is specified, it is ignored.</p> <p>Requires *CHANGE authority to an optical volume.</p> <p>For UDF formatted volumes, *X authority is required to each directory in the path name preceding the file and *W authority is required to the file.</p>	API not supported.
Close Directory (QHFCLODR)	None	API not supported.
Force Buffered Data (QHFFRCSF)	<p>If the volume media format is *UDF then data is forced to optical media.</p> <p>If the volume media format is not *UDF then data is forced to internal disk storage, not to optical media.</p> <p>For a file opened for read-only access, this API has no effect.</p>	API not supported.
Lock and Unlock Range in Stream File (QHFLULSF)	None	API not supported.

Table 5. Optical HFS API Restrictions (continued)

HFS APIs	Directly Attached Usage Notes	LAN-Attached Usage Notes
<p>Move Stream File (QHFMVSE)</p>	<p>If the source file is in the QOPT file system, *CHANGE authority is required to the optical source volume.</p> <p>If the target file is in the QOPT file system, *CHANGE authority is required to the optical target volume.</p> <p>When moving between the QOPT and QDLS file systems, file attributes are optionally copied depending on the global optical attribute CPYATR. This attribute can be displayed or changed using the CHGOPTA command.</p> <p>If the source file is on a UDF formatted volume, *WX authority is required to the parent directory and *X authority is required to each directory in the path name preceding the parent directory. *RW authority is required to the file.</p> <p>If the target file is on a UDF formatted volume, *WX authority is required to the parent directory and *X authority is required to each directory in the path name preceding the file.</p>	<p>API not supported.</p>
<p>Open Directory (QHFOPNDR)</p>	<p>Opening file system root (/QOPT) will allow both directly attached and LAN attached volumes to be returned on Read Directory Entries.</p> <p>Lock mode is ignored when opening the file system root.</p> <p>Lock mode of no lock is not supported. If requested, a lock mode of deny none is substituted.</p> <p>Requires *USE authority to the optical volume.</p> <p>For UDF formatted volumes, *X authority is required to each directory in the path name preceding the directory being opened and *R authority is required to the directory being opened.</p>	<p>API not supported.</p>

Table 5. Optical HFS API Restrictions (continued)

HFS APIs	Directly Attached Usage Notes	LAN-Attached Usage Notes
Read Directory Entries (QHFRDDR)	<p>QNAME is returned without the QOPT file system name.</p> <p>QNAME is the only field which is set for a LAN attached volume.</p> <p>QWRDTTM will always equal QCRTDTM.</p> <p>For files and directories, QACCDTTM will always equal QCRTDTM.</p> <p>For volumes, QACCDTTM will equal the last volume reference date.</p>	API not supported.
Rename Stream File (QHFRNMSF)	<p>Requires *CHANGE authority to the optical volume.</p> <p>For UDF formatted volumes, *WX authority is required to the parent directory and *X authority is required to each directory in the path name preceding the parent directory. *W authority is required to the file being renamed.</p>	API not supported.
Rename Directory (QHFRNMDR)	API not supported.	API not supported.

Control File System Functions to Optical Support

Optical support provides Control File System (QHCTLFS) functions to perform unique operations for the optical file system. These are optical specific functions that are not otherwise available through the HFS APIs. Different functions are available for directly-attached and LAN-attached optical devices.

The following control file system functions are available for directly-attached media libraries:

- SAV - saves a held optical file.
- RLS - releases a held optical file.
- SRD/VOL - performs a sector read to an optical volume.
- SRD/DEV - performs a sector read to an optical device.
- RTV/VOL - returns volume-specific information.
- GET - reads file data directly from the media with minimal caching.

Control File System Functions for Directly-Attached Optical

The following functions are available for directly-attached optical devices.

Save Held Optical File Function: Use the Control File System program to save a held optical file. A process must be allowed read access to a held optical file to save it.

The following is the syntax for the input buffer for the QHCTLFS program:

'SAV' + '/' + held-file-path + '/' + destination-file-path

For example:

- Input Data Buffer: SAV/VOLUME1/DIRECTORY1/FILE1//VOLUME2/DIRECTORY2/FILE2

- Input Data Buffer Length: 54

This function is also available using an option on the Work with Held Optical File (WRKHLDOPTF) display. However, unlike the save option on the Work with Held Optical File (WRKHLDOPTF) display, the save held optical file function of the control file system API does not automatically release a held file after it is saved. Therefore, an explicit release held optical file request is needed afterwards. For an explanation of held optical files, see “Held Optical Files” on page 85.

Release Held Optical File Function: Use the Control File System program to clear the held status of a file and release the optical file system from its obligation to write to the optical disk. A process must be allowed read and write access to a held file in order to release it; this means that no locks may currently be imposed on the file by other active jobs.

The following is the syntax for the input buffer for the QHFCTLFS program:

```
'RLS' + '/' + held-file-path
```

For example:

- Input Data Buffer: RLS/VOLUME1/DIRECTORY1/FILE1
- Input Data Buffer Length: 28

This function is also available using an option on the Work Held Optical File (WRKHLDOPTF) display. For an explanation of held optical files, see “Held Optical Files” on page 85.

Sector Read Function: The Control File System program can be used to do a sector reading of optical media. The sector read function is useful if the application knows precisely where data is stored on optical media. Sector read functions can be accomplished without opening and closing files, and independently of all HFS APIs. Multiple sectors may be read at one time.

There are two variations of the input buffer for issuing the Control File System sector read function:

```
SRD/VOL/volume_name/starting sector/number of sectors  
SRD/DEV/device_name/starting sector/number of sectors
```

Both return the range of sectors requested by the user. Sectors can be requested from an optical volume or optical device. For example, if an application wanted to read five sectors of optical volume VOL01 beginning at sector 1000, the following is requested:

```
SRD/VOL/VOL01/1000/5
```

Note: DEV is valid for stand-alone CD and DVD devices.

Retrieve Volume Information Function: Use the Control File System program to retrieve information about a particular volume.

The following is the input buffer format for the QHFCTLFS program:

```
RTV/VOL/volume_name
```

The format of the information returned in the output buffer is identical to the output file structure for volume attributes (QAMODVA) as shown in Appendix D, “Output File Support Structures,” on page 143. The length of the output buffer must be large enough to contain the QAMODVA structure.

Get File Data: You can use the Control File System (QHFCTLFS) HFS API to read a block of data from a file directly into your output buffer. This function improves performance when reading an entire file sequentially or when reading large blocks of data. The optical file system will not copy or cache the data as it does through normal Open, Read, and Close Stream File HFS APIs. When doing random reads to a file, the Open, Read, and Close Steam File option may still provide the best performance.

The following restrictions apply when using this API:

- Align output buffer on a 512 byte boundary.
- File offset must be 0 or a multiple of 4096.
- Maximum-read size is 16,000 sectors (16,384,000 bytes).
- The HFS API requires Shared No Update (*SHRNUP) access to the file.
- Calling program must be in User (not System) state.
- The HFS API requires *use authority to the volume.

Here is the syntax for the input buffer for the QHFCTLFS program:

'GET' + '/' + entire path + '/' + bytes to read + '/' + file offset

The following example will read 15MB from FILE.XXX. Starting at the beginning of the file with (offset=0):

- Input Data Buffer: GET/VOL1/DIR1/SUBDIR1/FILE.XXX//15728640/0
- Input Data Buffer Length: 42

The number of bytes read is returned in the "Length of Data Returned" parameter. In the above example if FILE.XXX is only 50KB in size, 51200 will be returned in the field. Therefore, it is not necessary to know the file size prior to issuing this request. Likewise, if 15728640 is returned in the "Length of Data Returned" parameter, the file is at least 15MB in size. More reads may be necessary to retrieve all the data.

It is not required that the number of bytes to read is a multiple of 4096. Although if it is not, data may be read into the output buffer beyond the number of bytes requested. This is because the device does I/O in blocks of 4096 bytes. Therefore, reading data in multiples of 4096 bytes is advise to avoid this problem.

Errors from Control File System (GET): The table below shows some common application errors that may occur using this API.

Table 6. Common errors for GET API

Message	Error
OPT1812 with 6030 as unexpected return code	File offset is beyond the end of file.
OPT1812 with A950 as unexpected return code	Output buffer is not 512 byte aligned.
OPT1860	Bytes to read is greater than the buffer size.
OPT1812 with C060 as unexpected return code	Attempted to read more than 16,384,000 bytes.
OPT1812 with C061 as unexpected return code	File offset is not a 4096 multiple.
CPF1F48	Input buffer is invalid. Verify the syntax.

Control File System Functions for LAN-Attached Optical Devices

The following control file system functions are available for LAN-attached media libraries.

- UPD/LAN - performs a dynamic refresh of the LAN volume lists.
- UPD/VOL - returns volume-specific information.
- RTV/VOL - returns volume-specific information.
- RTV/DIR - returns subdirectory and file entries for a specified directory.

Update Volume Information: Use the Control File System program to retrieve information about a particular volume or to update the internal list of available volumes on a LAN.

The following is the input buffer format for the QHFCTLFS program:

UPD/VOL/volume_name

It performs the following:

- UPD/VOL/volume-name: Using this input buffer format returns the amount of free space on a volume, total user space, media type, and flip-side volume ID. The format is shown here:
 - Bytes (1-32): Opposite-side volume ID.
 - Bytes (33): Reserved.
 - Bytes (34-37): User free space on the volume. This is a four-byte binary field.
 - Bytes (38-41): Total free space on the volume. This consists of the user free space on the volume plus the reserved space on the volume. The reserved space on the volume is determined when setting the volume-full threshold for the volume. This is a four-byte binary field.
 - Bytes (42): Media type. This is a one-byte binary field that can have the following values.
 - 0 = Non-valid Media or 3431 Standalone Drive
 - 1 = Write Once Read Many (WORM) media
 - 2 = Rewriteable media
 - Bytes (43): Magnitude of free space on the volume. This is a one-byte binary field that can have the following values:
 - 0 = Space field is in number of bytes.
 - 1 = Space field is in number of kilobytes (1024).
 - 2 = Space field is in number of megabytes (1048576).
 - Bytes (44): Magnitude of Total Space on the Volume. This is a one-byte binary field that can have the following values:
 - 0 = Space field is in number of bytes.
 - 1 = Space field is in number of kilobytes (1024).
 - 2 = Space field is in number of megabytes (1048576).
- UPD/LAN: Using this input buffer format updates an internal list of available volumes on all activated servers. You can perform this function after adding or removing cartridges from dataservers.

Retrieve Volume Information Function: Use the Control File System program to retrieve information about a particular volume.

The following is the input buffer format for the QHFCTLFS program:

RTV/VOL/volume_name

The format of the information returned in the output buffer is identical to the output file structure for volume attributes (QAMODVA) as shown in Appendix D, “Output File Support Structures,” on page 143. The length of the output buffer must be large enough to contain the QAMODVA structure.

The system uses format QAMODVA for volumes in all optical device types. While the format is the same, not all fields contain a value for LAN volumes. Check the “Output File Structure for Volume Attributes” on page 143 for the differences between the device types.

Retrieve Directory Information Function: Use the Control File System program to retrieve a list of files and subdirectories for a particular directory.

The following is the input buffer for the QHFCTLFS program:

```
RTV/DIR/volume_name/directory_name
```

The directory information is returned in the output buffer in the following format:

- CBdirectoryBCBdirectoryBCBfilenameBCBfilenameBB, whereas the following are:
 - C
 - D = Directory Entry
 - F = File Name Entry
 - B = EBCDIC Blank
 - BB = Double EBCDIC Blanks to indicate end of string

The output buffer must be at least 31KB long.

Standard Attribute Definition and Restrictions—Directly-Attached

Directory entries for files and directories have information that is associated with them called **attributes**. Each attribute consists of a name and a value. Some attributes generate automatically when creating the directory or file. These attributes are called **standard attributes**. Standard attributes start with the letter Q for ease of identification. All file systems use standard attributes. Several receive unique interpretation by the optical file system. LAN-attached optical devices have a different interpretation of standard attributes than directly-attached optical devices. For specific information that regards standard attribute differences for LAN and directly-attached optical devices, see “Optical File System Implementation of HFS APIs” on page 60. Also, you can find the meaning of these attributes for each HFS API in “Optical File System Implementation of HFS APIs” on page 60. The following is a basic definition of the standard attributes and their meaning with respect to optical support.

QALCSIZE Attribute

As an output field, QALCSIZE is the number of bytes allocated on optical disk by the file. It will always be 0 for directories.

When the QALCSIZE attribute is specified on Open Stream File during a write request, the media is checked to see if there is enough space available to allocate the amount specified. If there is not enough space available on the optical volume, message CPF1F61, No free space available on media, is issued. For more information on using this attribute, see “Media Capacity and Volume Threshold” on page 82

QACDRTM Attribute

This attribute is not supported by the optical file system. It is always the same as the file creation date and time (QCRTDTM).

QCRTDTM Attribute

This attribute indicates the creation date of a file or directory.

QWRTDTM Attribute

This attribute indicates the last date and time that data was written to an optical file. It does not reflect the date and time that the file attributes were last written.

QFILATTR Attribute

Support of this attribute is only by directly-attached optical support devices. The optical interpretation of the file flags is as follows:

- Read-only file: OS/400 provides full support of this attribute through the optical file system. When setting this attribute to ON (1), you cannot delete or overwrite the file.
- Hidden file: OS/400 maintains this attribute for the user application to manage, but does not fully support it by the optical file system. When setting this attribute to ON (1), the optical file system does not recognize the file as hidden. User applications require no special access to files with this attribute on.
- System file: OS/400 maintains this attribute for the user application to manage; but does not fully support it by the optical file system. When setting this attribute to ON (1), the optical file system does not recognize the file as a system file. User applications require no special access to files with this attribute on.
- Changed file: OS/400 supports this attribute by the optical file system. It is automatically set on(1) when a file is created or written to. You can only set it off(0) by using the Change Directory Entry Attributes (QHFCGAT) API.

Special Attributes Unique to Optical File System

Attributes for files and directories that are not standard and therefore not recognized by HFS are referred to by HFS as **extended attributes**. They are usually defined by a business application, but some are recognized by the optical file system as having special meaning.

OPT.CHGATDTM Attribute

This optical attribute reflects the last date and time that the file attributes were written. It is returned to the user application as an extended attribute through the Retrieve Directory Entry Attributes (QHFRVAT) command.

QOPT.IOMETH Attribute

This is a special extended attribute to the optical file system. Provided supported is only by directly-attached optical support devices; it is ignored by LAN support. The system also ignores this attribute when the media format is Universal Disk Format.

When an extended attribute of this name is passed by the application as the attribute name field in the Attribute Information Table (AIT) during an open stream file request, the optical file system knows that a special method of I/O is being requested. The optical file system retrieves the special method of I/O from the attribute value field in the AIT.

Currently, there is only one special method of I/O supported by the optical file system: **expanding buffer I/O**. You can request this method of I/O when the attribute value field for the QOPT.IOMETH attribute contains the value (EXPNBUFF). The optical software recognizes this special extended attribute as a

requested I/O method, and not as a normal extended attribute. It is not hereafter associated with the file in any way, and does not appear when attributes for the file are retrieved. All read operations for the process use expanding buffer I/O until the file is closed. Methodology and restrictions for using expanding buffer I/O are listed here. In order to determine if expanding buffer I/O should be used, see “Expanding Buffer I/O through HFS—Tailoring Read Requests for Performance” on page 84.

An HFS attribute in an attribute information table consists of several fields. These fields and the values you specify when opening a file for expanding buffer I/O are summarized in Table 7.

Table 7. Expanding Buffer Attribute Definition

Field	Data Type (see note)	Value for EBIO
Attribute Name	CHAR(*)	QOPT.IOMETH
Attribute Value	CHAR(*)	EXPNBUFF
Length Attribute Name	BIN(4)	0000000B
Length Attribute Value	BIN(4)	00000008
Note:		
<ul style="list-style-type: none"> • CHAR(*) indicates a variable number of bytes of character information. • BIN(4) indicates 4 bytes of binary information. • All character fields should be set in uppercase. 		

In addition to the values for attribute fields in Table 7, two additional fields are required to build an attribute information table:

- The number of attributes defined in the table.
- The table offset to each attribute, in bytes.

The Open Stream File (QHFOFNSF) API requires 10 bytes of open information as input. When you attempt to open a file for expanding buffer I/O, the open information is subject to the following restrictions:

- The action to take if a file exists must be to open the file.
- The action to take if a file does not exist must be to return an error.
- The lock mode for the file must be Deny Write or Deny Read/Write (exclusive).
- The access mode for the file must be Read Only.

If there is an expanded buffer I/O attribute in the attribute information table and any of these restrictions are not observed, an OPT1133 message is issued, indicating which of the fields in the open information was passed in error.

IBM provides online information on the format of attributes, the Attribute Information Table, or the Open Stream File API. Refer to the APIs topic in the Programming category of the Information Center.

Expanding Buffer I/O HFS API Restrictions: In addition to the restrictions that were detailed when opening a file for expanded buffer I/O, you cannot use the following APIs after a file is opened for expanding buffer I/O:

- Write Stream File
- Set Stream File Size
- Lock or Unlock Range in Stream File

Copying File Attributes Using HFS

When you copy files, using the hierarchical file system, between QOPT and QDLS file systems, the target file is assigned either default file attributes or the file attributes of the source file. This depends on the value you specify for the copy attributes (CPYATR) global value on the Change Optical Attributes (CHGOPTA) command.

When the CPYATR global value is specified as *NO on the CHGOPTA command, default file attributes are created for files that are copied between the QOPT and QDLS file systems.

When the CPYATR global value is specified as *YES on the CHGOPTA command, file attributes from the source file are copied to the target file for copies between the QOPT and QDLS file system.

Copying Attributes from QDLS to QOPT

In copies or moves from QDLS to QOPT, the following default attributes are assigned to the target file:

- Standard file attributes:
 - Creation date and time is set to the current date and time
 - Modification date and time is set to the current date and time
 - Access date and time is set to the current date and time
 - The QFILATTR standard attribute is set to '00000': the file is not read-only; the file is not hidden; the file is not a system file; the file is not a directory; and the file has not changed since it was last archived or created.
- No DIA document attributes are copied.
- No user-defined extended attributes are copied.

The file name (QNAME) and file size (QFILSIZE) are maintained.

Copying Attributes from QOPT to QDLS

In copies or moves from QOPT to QDLS, the following default attributes are created:

- Standard file attributes:
 - Creation date and time is set to the current date and time
 - Modification date and time is set to the current date and time
 - Access date and time is set to the current date and time
 - The QFILATTR standard attribute is set to '00000': the file is not read-only; the file is not hidden; the file is not a system file; the file is not a directory; and the file has not changed since it was last archived or created.
- DIA document attributes:
 -
 - DIA.CA04C700 (text description) is set to the file name.
 - DIA.CA04C701 (profile GCID) is set to code page 697 and character set 500.
 - DIA.CA04C706 (file type) is set to X'000E' (PC file).
 - DIA.CA04C720 (library assigned document name) is assigned to represent this file.
 - DIA.CA04C708 (last changed date and time) is set to the current date and time.
 - DIA.CA04C707 (creation date and time) is set to the current date and time.

- DIA.CA04C710 (NLS information) is set to the language ID and country or region ID of the job.
- DIA.CA04C740 (file date and time) is set to the current date and time.
- No user-defined extended attributes are copied.

The file name (QNAME) and file size (QFILSIZE) are maintained.

Using the Integrated File System Support Interface to Optical Support

The integrated file system support provides a UNIX-type interface that you can use to maintain optical files and directories. LAN-attached optical devices do not support this interface .

The integrated file system support for optical support consists of:

- **UNIX-type APIs:** These APIs are C language functions that can be used in ILE C for OS/400 programs.
- **Generic commands:** These are system CL commands that allow interface to optical support.
 - **TCP/IP File Server Support for OS/400:** This licensed program offering allows PC users to access the files and directories of optical systems or other file systems. For more information, see the *OS/400 Server Concepts and Administration* book.

IBM provides online information about using the HFS or UNIX-type APIs, and on the generic commands. Refer to the APIs topic in the Programming category of the Information Center.

Optical File System Implementation of Integrated File System Support

Like all file systems, the optical file system has unique rules and restrictions for applications that access optical functions through the integrated file system. Several of the UNIX-type APIs and generic commands are not supported. Others are only partially supported, or restricted. Table 8 on page 76 and Table 9 on page 81 summarize these differences for the optical file system.

Table 8. Optical Implementation of UNIX-type APIs

UNIX-type API	Supported	Comments and Usage Notes
access (Determine File Accessibility)	Yes	<p>Requires *X authority to the parent optical volume. For non-UDF volumes, no other authority is required. For UDF formatted volumes the following authorization rules apply:</p> <ul style="list-style-type: none"> • Requires *X authority to each directory in the path preceding the object tested. • Requires *R authority when R_OK is specified. • Requires *W authority when W_OK is specified. • Requires *X authority when X_OK is specified. • Requires *RX authority when R_OK X_OK is specified. • Requires *WX authority when W_OK X_OK is specified. • Requires *RX authority when R_OK W_OK is specified. • Requires no authority when F_OK is specified.
accessx (Determine File accessibility based on the who parameter)	Yes	<p>Does not require *X authority to the parent optical volume. For UDF volumes the following authorization rules apply:</p> <p>Valid values for who are:</p> <ul style="list-style-type: none"> • ACC_INVOKER • ACC_SELF • ACC_ALL • ACC_OTHERS <ol style="list-style-type: none"> 1. Requires *R authority when R_OK is specified 2. Requires *W authority when W_OK is specified 3. Requires *X authority when X_OK is specified <p>Authority checks are mutually exclusive.</p>
chdir (Change Current Directory)	Yes	<p>Requires *X authority to the parent optical volume.</p> <p>For non-UDF volumes, no other authority is required.</p> <p>For UDF formatted volumes, *X authority is required to each directory in the path</p>
chmod (Change File Authorizations)	Yes	<p>Only supported for UDF formatted optical volumes. Requires *CHANGE authority to the parent optical volume. Requires *X authority to each directory in the path preceding the object. To perform this operation you must be the owner of the file or have *ALLOBJ special authority.</p>
chown (Change Owner and Group of File)	Yes	<p>Only supported for UDF formatted optical volumes. Requires *CHANGE authority to the parent optical volume. Requires *X authority to each directory in the path preceding the object. To perform this operation, you must be the owner of the file, or have *ALLOBJ special authority. Files and directories on non-UDF formatted volumes are owned by QDFTOWN user profile.</p>

Table 8. Optical Implementation of UNIX-type APIs (continued)

UNIX-type API	Supported	Comments and Usage Notes
close (Close File Descriptor)	Yes	
closedir (Close Directory)	Yes	
creat (Create or Rewrite File)	Yes	Requires *CHANGE authority to the parent optical volume. For non-UDF volumes, no other authority is required. For UDF formatted volumes, *X authority is required to each directory in the path and *WX authority to the parent directory. The change and modification time stamps for the parent directory are not updated.
dup (Duplicate Open File Descriptor)	Yes	
dup2 (Duplicate Open File Descriptor to Another Descriptor)	Yes	
fchmod (Change File Authorizations by Descriptor)	Yes	Only supported for UDF formatted optical volumes. To perform this operation you must be the owner of the file or have *ALLOBJ special authority.
fchown (Change Owner and Group of File by Descriptor)	Yes	Only supported for UDF formatted optical volumes. To perform this operation you must be the owner of the file or have *ALLOBJ special authority. Files and directories on non-UDF formatted volumes are owned by QDFTOWN user profile.
fcntl (Perform File Control Command)	No	
fpathconf (Get Configurable Path Name Variables by Descriptor)	Yes	
fstat (Get File Information by Descriptor)	Yes	Owner, group, and other mode bits are always on, regardless of the user's authority to the file. File access time stamp is not changed.
fsync (Synchronize Changes to File)	Yes	For UDF formatted volumes, data is forced to optical disk. For non-UDF formatted volumes, data is forced to internal disk storage recoverable through held optical files.
ftruncate (Truncate File)	Yes	
getcwd (Get Current Directory)	Yes	Requires *X authority to the parent optical volume. For non-UDF volumes, no other authority is required. For UDF formatted volumes, *RX authority is required to each directory in the path name preceding the object.
getegid	Yes	
geteuid	Yes	
getgid	Yes	
getgrid	Yes	
getgrnam	Yes	

Table 8. Optical Implementation of UNIX-type APIs (continued)

UNIX-type API	Supported	Comments and Usage Notes
getgroups	Yes	
getpwnam	Yes	
getpwuid	Yes	
getuid	Yes	
ioctl (Perform File I/O Control Request)	No	
link (Create Link to File)	No	QOPT does not support links.
lseek (Get File Read/Write Offset)	Yes	
lstat (Get File or Link Information)	Yes	File access time stamp is not changed. Requires *X authority to the parent optical volume. For non-UDF volumes, no other authority is required. For UDF formatted volumes, *X authority is required to each directory in the path preceding the object and *R authority to the object.
mkdir (Make Directory)	Yes	Requires *CHANGE authority to the parent optical volume. For non-UDF volumes, no other authority is required. For UDF formatted volumes, *X authority is required to each directory in the path and *WX authority to the parent directory. The change and modification time stamps for the parent directory are not updated. Owner ID and group ID are not set.
open (Open File)	Yes	If opening file for write, *CHANGE authority is required to parent optical volume. If opening a file for read, *USE authority is required to the parent optical volume. For UDF formatted volumes, the following additional authorization rules apply: <ul style="list-style-type: none"> • Requires *R authority when object is being opened O_RDONLY. • Requires *W authority when object is being opened O_WRONLY. • Requires *RW authority when object is being opened O_RDWR. • Requires *WX to the parent directory when object does not exist and O_CREAT is specified.
opendir (Open Directory)	Yes	Requires *USE authority to the parent optical volume. For UDF formatted volumes, *X authority is required to each directory in the path preceding the object, and *R authority to the object being opened.
pathconf (Get Configuration Path Name Variables)	Yes	

Table 8. Optical Implementation of UNIX-type APIs (continued)

UNIX-type API	Supported	Comments and Usage Notes
Qp0lGetPathFromFileId	Yes	
Qp0lRenameKeep	Partial	<p>QOPT does not support renaming a directory. Object must be a file.</p> <p>Requires *CHANGE authority to the parent optical volume. For non-UDF volumes, no other authority is required. UDF formatted volumes require *X authority is required to each directory in the path, and *WX authority to the parent directory, and *W authority to the file. If renaming the volume, *RWX is required to the root (/) directory of the volume.</p> <p>New and old files must exist in the same directory.</p>
Qp0lRenameUnLink	Partial	<p>QOPT does not support renaming a directory. Object must be a file.</p> <p>Requires *CHANGE authority to the parent optical volume. For non-UDF volumes, no other authority is required. UDF formatted volumes require *X authority is required to each directory in the path, *WX authority to the parent directory, and *W authority to the file. If renaming the volume, *RWX is required to the root (/) directory of the volume.</p> <p>Object that is identified by a new path cannot exist.</p>
read (Read from File)	Yes	The file access time is not updated. When reading from files on volumes formatted in Universal Disk Format (UDF), byte locks on the range being read are ignored. The same is true for readv().
readdir (Read Directory Entry)	Yes	The directory access time is not updated.
readlink (Read Value of Symbolic Link)	No	QOPT does not have symbolic links.
rename (Rename File or Directory)	Partial	<p>QOPT does not support renaming a directory. Object must be a file or a volume.</p> <p>Requires *CHANGE authority to the parent optical volume. For non-UDF volumes, no other authority is required. UDF formatted volumes require *X authority is required to each directory in the path, *WX authority to the parent directory, and *W authority to the file. If renaming the volume, *RWX is required to the root (/) directory of the volume.</p> <p>Object that is identified by a new path cannot exist.</p>
rewinddir	Yes	

Table 8. Optical Implementation of UNIX-type APIs (continued)

UNIX-type API	Supported	Comments and Usage Notes
rmdir (Remove Directory)	Yes	Requires *CHANGE authority to the parent optical volume. For non-UDF volumes, no other authority is required. For UDF formatted volumes, *X authority is required to each directory in the path and *WX authority to the parent directory. Change and modification time stamps for the parent directory are not updated. Operation will not be allowed if the directory is busy.
stat (Get File Information)	Yes	File access time stamp is not changed. Requires *X authority to the parent optical volume. For non-UDF volumes, no other authority is required. For UDF formatted volumes, *X authority is required to each directory in the path preceding the object and *R authority to the object. When issued to an optical volume, the size returned is the volume capacity or 2,147,483,647, whichever is smaller.
symlink (Make Symbolic Link)	No	QOPT does not support symbolic links.
sysconf (Get System Configuration Variables)	Yes	
unmask (Set Authorization Mask for Job)	Yes	
unlink (Remove Link to File)	Yes	Requires *CHANGE authority to the parent optical volume. For non-UDF volumes, no other authority is required. For UDF formatted volumes, *X authority is required to each directory in the path and *RX authority to the parent directory. Change and modification time stamps for parent directory are not updated. Link to a file cannot be removed when a job has the file opened.
utime (Set File Access and Modification Times)	No	QOPT does not support setting the file access or modification time.
write (Write to File)	Yes	Change and modification time stamps for the file are updated when the file is closed. When writing to files on volumes formatted in Universal Disk Format (UDF), byte locks on the range being written are ignored. The same is true for writev().

Refer to the iSeries Security Reference book for authorities that are required to issue generic commands.

Table 9. Optical Implementation of Generic Commands

Generic Command	Supported	Comments and Restrictions
ADDLNK	No	
CHGAUD	No	
CHGAUT	Yes	Supported only for UDF formatted optical volumes. ¹
CHGCURDIR	Yes	
CHGOWN	Yes	Supported only for UDF formatted optical volumes.
CHGPGP	Yes	Supported only for UDF formatted optical volumes.
CHKIN	No	
CHKOUT	No	
CPY	Yes	
CRTDIR	Yes	Command will fail if attempt is to create /QOPT or next level directory, which represents a volume.
DSPAUT	Yes	
DSPCURDIR	Yes	
DSPLNK	Yes	
ENDJRN	No	
MOV	Partial	QOPT does not support moving a directory, if it contains files or subdirectories. QOPT does not support moving a volume.
RMVDIR	Partial	QOPT does not support RMVLNK(*YES).
RMVLNK	Yes	
RNM	Partial	QOPT does not support renaming a directory.
RST	Partial	QOPT support restoring an entire volume using SUBTREE (*STG).
RTVCURDIR	Yes	
SAV	Partial	QOPT supports saving an entire volume using SUBTREE (*STG).
SAVRST	No	
STRJRN	No	
WRKAUT	Yes	Supported only for UDF formatted optical volumes. ^{1, 2}
WRKLNK	Yes	

Notes:

- To perform this operation you must be the owner of the file or have *ALLOBJ special authority.
QOPT does not maintain or honor object level authorities associated with optical files and directories. Therefore, any attempt to change or revoke object

level authorities is not allowed. The only allowed value for the New object authorities (OBJAUT) parameter is *SAME.

Specifying *EXCLUDE for the New data authorities (DTAAUT) parameter is not allowed. Command parameter rules require that if *EXCLUDE is specified for the New data authorities parameter a value of *NONE must be specified for the New object authorities parameter. See restriction 'b' above for an explanation.

If the desire is to revoke authority associated with the owner, group or other user *NONE may be specified as a value for the New data authorities parameter. In this case the specified user and their data authorities are removed from the list of authorized users.

QOPT does not maintain or honor a private authority list. An attempt to assign New data authorities to a user other than the owner, group or other (*PUBLIC) is not allowed.

2. QOPT does not maintain or honor a private authority list. An attempt to add a new user, option 1 from the WRKAUT display, and assign new data authorities to a user other than the owner, group or other (*PUBLIC) is not allowed.

Option 4 is not supported to remove user from list of authorized users. Select and prompt (F4) option 2 on the user you wish to remove. The New data authorities parameter (DTAAUT) must be set to *NONE and the New object authorities parameter (OBJAUT) set to (*SAME). See 1 on page 81.

Thread Safety

In Version 4 Release 4 of OS/400, the optical file system (QOPT) is thread-safe through the integrated file system application program interfaces. You can spawn multiple threads in the same process that is reading and writing files in QOPT. This holds true for all optical volumes that are accessed using the integrated file system APIs. This includes all optical volumes except those in LAN-attached optical media library devices.

Tips and Techniques for the Application Programmer

This topic describes how the optical file system manages file data so application programmers can optimize their applications. Since applications have different requirements, this topic does not suggest the best way to write an optical application. It does, however, provide explanations that all application programmers could find useful.

You should use this topic to determine the best way to handle optical file management, either through the HFS or UNIX-type APIs. You should use this topic only for applications to directly-attached optical support. **Concepts in this topic do not apply to optical LAN support.**

IBM provides online information on the HFS or UNIX-type API functions. Refer to the API topic in the Programming category of the Information Center.

Media Capacity and Volume Threshold

One thing to consider when writing to optical media is the possibility of reaching the media capacity or threshold. The optical file system provides a logical threshold capability to help applications protect themselves from reaching the absolute volume capacity. Definition of the logical threshold occurs when the volume initializes, and is unique for each volume. You can change this threshold by using the Change Optical Volume command.

Note: The logical volume threshold is only applicable for high performance optical file system (HPOFS) media format. For UDF media format, the logical volume threshold is always 100% and cannot be changed.

You should devise a strategy to deal with the situation when the media becomes full. This is especially true when writing to WORM media. You might consider the following questions:

- How should I use the volume threshold?
- What should I do when the volume is full?
- How can I prepare for a volume-full condition?

The volume threshold is provided to allow applications to prepare for an actual volume-full condition. When WORM media becomes full, there can be no further write operations. Depending on the requirements of the application, the threshold can be used in various ways to prepare for the media becoming physically full.

For example, an application might write groups of spooled files to optical disk. After each group is written, an additional file might be written that contains an index to the spooled files just written. Without the index, the spooled files could be useless. Unless the application can manage the media capacity, the volume might run out of space before the index file can be written. One way around this is to set the volume threshold to 99%. When a No space available message is issued, the application can then increase the threshold to 100% and write any necessary additional files.

Managing Media Capacity on a Per-File Basis

An application might also need to manage the media capacity on a per-file basis. Following are a few methods you can use to decide if a file will fit on the media:

1. Handle error on a close operation

Assume an optical volume is initialized to a 95% threshold and an application writes files until the volume threshold is reached. When the threshold is reached, the application will receive message CPF1F61, No free space available on media.. At this point, the volume threshold can be increased to 97% (or anything else up to 100%) by using the CHGOPTVOL command. You can then attempt to close the file.

2. Specify QALCSIZE on the Open Stream File HFS API

Another method to determine if a file will fit on a volume is by specifying an allocation size (QALCSIZE) on an open stream file. On an open stream file, the system can pass a value in attribute QALCSIZE. This attribute is valid when the open operation is for create or replace; otherwise, it is ignored. Specifying a value for QALCSIZE results in comparing the specified value against the space available on the volume. If the space available is less than QALCSIZE, then the system issues message CPF1F61. Having the space available exceed the QALCSIZE allows the open operation. Only on the first open instance of a file honors this attribute. If specified by more than one opening of a file, the system ignores the additional attributes.

Note: This does not actually allocate space on the optical volume at the time of the open operation. It simply checks the volume to see if the number of bytes that are requested are available.

There are drawbacks to using this method:

- a. You need to know at the time the open request is made the size of the file you are creating.

- b. If multiple jobs are writing to the same media, there is no guarantee that by the time the data is written, the space will still be available.

If the size of the file is known prior to the time the open request is made, and there will not be other jobs writing to that volume during the time your file is open, this is an excellent method to check media capacity before creating a file.

3. Retrieve space available on a volume

Another method is to have the application retrieve the space available on the volume. You can do this by using the Display Optical (DSPOPT) command through output file support. The output file can then be read to retrieve the number of bytes assumed to be available on the media.

Expanding Buffer I/O through HFS—Tailoring Read Requests for Performance

An alternative method of opening a stream file through HFS can improve performance for applications that typically read portions, but not all, of the data in large optical files. This alternate method of input/output is referred to as **expanding buffer I/O**. Expanding buffer I/O is available only to HFS API applications when accessing HPOFS or ISO9660 formatted media. This attribute is ignored when the media format is UDF.

Note: Using the HFS APIs, optical file data is buffered into a virtual optical file in OS/400 main storage. If expanding buffer I/O is not selected as an option, the size of this buffer is equal to the size of the actual optical file. For example, a 100MB file on optical media has a 100MB buffer when the file is opened through the HFS API, Open Stream File. The performance cost for overhead operations involving the optical buffer is proportional to the buffer size. The time it would take to read one byte of a 100MB file is substantially greater than reading one byte of a 50KB file.

When an optical file is opened for expanding buffer I/O, the size of the buffer begins at zero and expands as data is read into the buffer as requested by the application. The minimum amount of the size expansion is 256KB. The buffer expands only if the requested data is contained within a logical 256KB page that is not yet contained in the buffer. For these reasons, the amount of time it would take to read one byte of a 100MB file opened for expanding buffer I/O should be roughly identical to the time to read one byte of a 50KB file opened in the same manner.

Situations in Which Expanding Buffer I/O is Especially Useful

Expanding buffer I/O should be considered as an option for improving the performance of reads if any of the following is true.

- The typical size of an optical file to be read is greater than 256KB.
- The amount of data read from the optical file between the open and close stream file is a fraction of the total file data. The exact fraction would be impossible to specify, but the performance improvements that are achieved will be greater the smaller the fraction. For example, an application that used expanding buffer I/O to read 25KB of a 50MB file would experience much greater performance improvements than an application that read 45MB of the same file. An application that reads the entire 50MB example file 40KB at a time through multiple reads probably would not experience any performance improvement using expanding buffer I/O.

- The application will not issue the Set Stream File Size, Lock-Unlock Byte Range, or Write Stream file APIs while the file is open for expanding buffer I/O. For further expanding buffer restrictions, see “Expanding Buffer I/O HFS API Restrictions” on page 73.

For specific details on how to implement expanding buffer I/O, see “Special Attributes Unique to Optical File System” on page 72.

Force Buffered Data or fsync() APIs

When creating or updating optical files, the data is not guaranteed to exist on optical disk until the file is successfully closed. Optical file data can, however, be synchronously written to nonvolatile storage using either the HFS API Force Buffered Data (QHFFRCSF) or the fsync() UNIX-type API. The type of nonvolatile storage is different depending on the optical media format.

For High Performance Optical File System (HPOFS), all file data will be written to the internal disk storage. The data can then be recovered through a Held Optical File if a power loss or other unexpected error occurred which prevented the file from being closed.

For Universal Disk Format (UDF), all file data will be written to optical disk when a force is issued. No recovery is required if a power loss or other unexpected error occurred which prevented the file from being closed. However, if writes were issued after the data was forced and the close was never successful, the file data is unpredictable since the writes that followed the force are asynchronous and the data may or may not have been written to optical.

Held Optical Files

Held optical files are virtual files that were never successfully written to optical media. A virtual file becomes *held* if an error occurs during the close operation of a file on a non-UDF formatted volume. You can manage these virtual files by using application interfaces and optical utilities. No creation of held files occurs for files that fail to archive on UDF formatted volumes.

Using the example cited in “Media Capacity and Volume Threshold” on page 82, assume an application reaches the volume threshold when writing to a file. This time, however, the absolute volume capacity is reached. The file is too large to fit on the volume. Because increasing the volume threshold will not help, another solution is needed. When the close request fails, the virtual file becomes held. Using the Work With Held Optical Files command, this virtual file can be saved to another volume. If desired, the file can be saved under a different name. The save request can also be performed using a control file system function. For more information on the optical control file system support, see “Control File System Functions to Optical Support” on page 67.

Volume, Directory, and File Considerations

The following topics discuss considerations for working with optical volumes, directories, and files.

Working with Volumes

Consider the following terms when referring to volumes:

Online

The volume is mounted in a drive under the read/write heads.

Near Online

The volume is in the optical media library, but not online. The volume may be in a storage slot or the opposite side of an online volume.

Removed

The volume is not physically in an optical media library, but volume information for the volume was kept when the volume was removed.

Offline

The volume is in an optical device, but the device is either powered off, varied off, or no longer connected.

Consider the following characteristics of optical volumes:

- An optical volume is one side of an optical cartridge.
- An optical cartridge contains two volumes.
- All volume names must be unique.
- Depending on the optical media density and type, the capacity of a volume can range from a few hundred megabytes to many gigabytes.
- Normally, a near online volume takes less than ten seconds to become an online volume. This requires the volume to be mounted into a drive.
- The number of drives in the optical media library determines how many volumes can be online at any time. Only one volume can be mounted in a drive (online) at one time. The remaining volumes in the library are near online.
- Volumes are generally independent of each other, with one exception. The two volumes on the same cartridge can never be completely independent. Both volumes on a cartridge can never be online at the same time. Copying between two volumes on the same cartridge can be done, but it requires the cartridge to be “flipped” several times to copy all of the requested files.
- There is no limit to the number of removed volumes that can exist.

How an application manages volumes depends almost entirely on the requirements of the application. Data should be written to volumes strategically, depending on the desired retrieval time in the future. If it is not desirable to wait for a near online volume to become online, the application might need to be set up so the most likely volumes to be accessed are online.

Working with Directories

The only limit to the number of directories that can be created on a volume is the capacity of the media. This restriction also applies to the number of files that can exist in an optical directory. Directories are not required to exist for files to be stored on a volume. If desired, all files can be stored in the root directory of a volume. The root directory is the “/” directory that is created when a volume is initialized. This root is not considered a directory in the traditional sense since it cannot be created or deleted like other directories. The root directory will always exist on initialized optical volumes.

Directories can be used to categorize optical files into more manageable subsets. Directories can contain files from a particular time period, subject, characteristic, or any combination of these. For example, there may be a directory SPOOLFILES with subdirectories YEAR_1994 and YEAR_1995. Taking this one step further, there could be subdirectories within these subdirectories named MONTH_MARCH and MONTH_APRIL. This structure could be represented as follows:

```
/SPOOLFILES  
/YEAR_1994  
/MONTH_MARCH
```

```
/MONTH_APRIL
/YEAR_1995
/MONTH_MARCH
/MONTH_APRIL
```

Fully qualified directory names for this example look like this:

```
/SPOOLFILES
/SPOOLFILES/YEAR_1994
/SPOOLFILES/YEAR_1994/MONTH_MARCH
/SPOOLFILES/YEAR_1994/MONTH_APRIL
/SPOOLFILES/YEAR_1995
/SPOOLFILES/YEAR_1995/MONTH_MARCH
/SPOOLFILES/YEAR_1995/MONTH_APRIL
```

Directories can be useful when categorizing files, but they are not necessary. Like volume names, directory names must be unique within the same volume. For example, volume VOL001 cannot have two directories named DIR001. Volume VOL001 can, however, have a DIR001 directory and a DIR000/DIR001 directory. Also, a DIR001 directory can exist on volume VOL001 and volume VOL002. For information on directory naming conventions, see “Path Names.”

Working with Files

The size of optical files depends almost entirely on the requirements of the application and the users of those files. The size of an optical file (accessible through HFS or the integrated file system) can range from 0 bytes to 4,294,705,152 bytes depending on the capacity of a volume. The physical size of the target piece of media is limited by the amount of free space available.

When selecting optimal file sizes for your application, pay special attention to the following considerations:

- The amount of system disk unit or main storage on the iSeries server
- How the data will be read (sequentially or randomly)
- Whether the entire file will usually be retrieved, or just a small portion
- Whether files will be updated once they are written to the volume

Generally, the larger the file, the better the performance and media use. When larger files are used, less media space is taken up by file directory information and more is used for actual data. Also, the performance related to file size is not a linear comparison. It does not take twice as long to write 20KB of data as it does to write 10KB of data. Performance (KB/second) improves as the amount of data being read or written increases.

Path Names

The term **path** refers to a file-system name, volume name, directory name, and file name.

Path Names for Volumes in Directly-Attached Devices

The following example illustrates the format for a path name on a directly-attached device. The forward slash (/) is used as a separator character. The path name must begin with a forward slash and contain no more than 294 characters.

```
/QOPT/VOL_NAME/DIRECTORY_NAME/SUB_DIR1/.../SUB_DIRn/FILE_NAME
```

QOPT refers to the optical file system. You must use it to qualify the optical file system when issuing calls to optical support through the HFS API or the Unix-type APIs. The portion of the path following the file system name cannot contain more than 289 characters. Following are rules for using path names:

- A path name can consist of any EBCDIC characters, except the characters that are listed below:
 - X'00' through X'3F'.
 - X'FF'.
 - The quotation mark (").
 - The asterisk (*).
 - The less than (<) and greater than (>) signs.
 - The question mark (?).
 - The hyphen (-).
 - The back slash (\).

When accessing UDF formatted volumes through the integrated file system APIs, the only characters not valid are X'00' through X'3F', X'FF', and back slash.

- The volume identifier can be a maximum of 32 characters for HPOFS media format, and a maximum of 30 characters for UDF media format. The identifier must contain only alphabetic characters (A through Z), numeric characters (0 through 9), a hyphen (-), an underscore(_), or a period (.). The first character must be alphabetic or numeric, and the identifier cannot contain blanks.
- You can include one or more directories in the path name, but it is not required. The total number of characters in all of the subdirectories together cannot exceed 256
- The file name is the last element in the path. The directory length in the path limits the file name length. The directory name and file name combined cannot exceed 256 characters. The preceding forward slash of the directory name is considered part of this 256 characters.

Path Names for Volumes in LAN-Attached Devices

The following example illustrates the format for a path name on an optical volume in a LAN-attached optical device. The forward slash (/) is used as a separator character. The path name must begin with a forward slash and contain no more than 261 characters.

```
/QOPT/VOL_NAME/DIRECTORY_NAME/SUB_DIR1/.../SUB_DIRn/FILE_NAME
```

QOPT refers to the optical file system, and must be used to qualify the optical file system when issuing calls to optical support through the HFS or integrated file system APIs. The portion of the path following the file system name cannot contain more than 256 characters. Following are rules for using path names on LAN-attached devices:

- See the *IBM 3995 LAN Optical Library Dataserver* book for the allowed characters for path names.
- The volume name is required and can contain a maximum of 32 characters.
- One or more directories can be included in the path name, but it is not required. The total number of characters in all of the subdirectories together cannot exceed 254.
- The file name is the last element in the path. The file name length is limited by the volume and directory length in the path. The volume name, directory name, and file name combined cannot exceed 256 characters. The preceding forward slashes of the volume and directory name are considered part of the 256 characters.

Chapter 8. Optical Volume Backup

This chapter describes the methods that you can use to back up optical data from one optical volume to another optical volume or tape.

Defining Your Backup Strategy

There is no one perfect backup strategy that meets everyone's needs. Therefore, it is important to define your backup requirements before you decide on a backup strategy. Use the following questions to help you determine your backup strategy:

- Do I need backups?
 - Could information be re-created easily?
 - If I do not have backups, how will that affect my business?
 - Am I legally required to have backups?
- How frequently should my backups be done?
 - Daily
 - Weekly
 - Monthly
 - Only when a volume is full
- How will backups be done?
 - Incremental backup
 - Partial or selective backups
 - Complete backups
- When do I want the system to perform the backups?
 - During first, second, or third shift
 - On the weekend
 - Will there be other contentions for the optical drives?
- Will target volume contain backups for one or multiple volumes?
- How long do I retain source information after a backup has completed?
- What type of availability is needed for volumes?
 - In optical media library
 - Out of optical media library, but on-site
 - Out of optical media library, and off-site

This is not a complete list of items to consider when deciding on a backup strategy but rather a foundation on which you can build.

Your Backup Options

Five options are available to support backing up your optical data. See the specific command for a discussion on comparative performance and suggested use.

1. Duplicate Optical (DUPOPT) CL command (QOPT specific command).
2. Save (SAV) CL command and Restore (RST) CL command (Integrated File System general purpose commands).
3. Copy Optical (CPYOPT) CL command to volume type *PRIMARY (QOPT specific command).

4. Copy Optical (CPYOPT) CL command to volume type *BACKUP
5. Copy Object (CPY) CL command (Integrated File System general purpose command).

IBM provides online information about the CPY, SAV, and RST CL commands. Refer to the *CL* section that is found in the *Programming* category of the **Information Center** at the following Web site - <http://publib.boulder.ibm.com/pubs/html/as400/infocenter.htm>.

DUPOPT

The Duplicate Optical (DUPOPT) command creates a duplicate optical volume. The created duplicate volume is identical to the original volume except for the volume identifier and the time created.

Performance improves significantly when you use DUPOPT instead of CPYOPT to back up an entire volume from scratch.

The DUPOPT command has the following requirements:

- Two optical drives are required.
- The source and target volumes must have the same sector size.
- The source and target volumes can not be on opposite sides of the same cartridge.
- The source and target device types must be the same type (for example, optical library to optical library or optical stand-alone device to optical stand-alone device).
- If the source media type is *WORM, the target media type can be either *WORM or *ERASE.
- If the source media type is *ERASE, the target media type must be *ERASE.
- If the source media type is *DVD-RAM, the target media type must be *DVD-RAM.
- If the target media type is *WORM, it must be uninitialized.
- If the target media type is *ERASE, all data currently on the target volume will be lost.
- Once the DUPOPT command begins operation, the system will not interrupt the process. The system will not schedule any other work in the drives until the command completes.

Use the DUPOPT command when you want to copy the entire volume or for incremental backups of your source volume. For an entire backup, wait until your source volume is full before you use the DUPOPT CL command.

DUPOPT always makes a complete copy of your source volume. However, you could use it to make incremental backups of your optical volumes in the following manner:

1. Determine how often you wish to backup your source and how many backup copies you wish to keep.
2. Use DUPOPT to duplicate your source media to a target media that has a media type of *ERASE. This will give you an exact copy of the source media.
3. Continue making duplicates of your source volume as often as you wish until your source volume becomes full.

4. Once the source media is full and you have a successful final copy, you can reuse all previous target media for backups of other source media.
5. If your source media type is *WORM, prior to your final backup determine whether or not your final target media type needs to be media type *WORM or *ERASE.

Attention: If the DUPOPT command does not complete successfully or it ends for any reason while processing, the backup is unsuccessful. In addition, if the target media type is *WORM, the target volume may no longer be usable.

Enhancements

Several enhancements have been made since the introduction of DUPOPT command:

- DUPOPT no longer requires the source and target volume to be in the same optical library.
- For target media with the *HPOFS media format, the target media capacity no longer has to be identical to the source media capacity. It can now be equal to or larger than the source media.
- For *WORM media the target media type no longer has to be *WORM. You can use either *WORM or *ERASE media.
- Improved device error recovery.
- Improved performance.
- The addition of the cross device support (library to library, stand-alone to stand-alone).
- The addition of the unload support for stand-alone devices.
- When duplicating a volume in a stand-alone or virtual device, you can create an exact copy with the same volume identifier as the source.

In the following example, the system duplicates the volume in optical device OPT01 to the volume in optical device OPT02. Upon completion, the target volume identifier will be MYBACKUP, and the system unloads the media from the device.

Note: The system only supports the unload option for stand-alone devices.

```
> DUPOPT FROMVOL(*MOUNTED) TOVOL(*MOUNTED) NEWVOL(MYBACKUP)
  FROMDEV(OPT01) TODEV(OPT02) FROMENDOPT(*LEAVE) TOENDOPT(*UNLOAD)
```

CPYOPT

This section discusses the Copy Optical (CPYOPT) CL command and some of the optional parameters which you can specify. Use the CPYOPT command to copy optical files and directories between optical volumes. You can use CPYOPT to copy any of the following:

- All files and directories for an entire volume.
- All files and subdirectories for a directory.
- All files for a directory.
- A single file.

There are two optical volumes types, and they are: *PRIMARY and *BACKUP. *PRIMARY is the normal volume type which can be written to by user applications. *BACKUP is a special volume type which only can be written to by special optical commands. Later, this chapter discusses the *BACKUP volume type in more detail. You can use CPYOPT to copy files between the following volume

types:

FROM Volume
*PRIMARY
*PRIMARY
*BACKUP

TO Volume
*PRIMARY
*BACKUP
*PRIMARY

Key parameters

There are several parameters which you can specify in order to help you select the files that you want copied, and they include:

- Select files to copy (SLTFILE)
- Copy subdirectories (CPYSUBDIR)
- Starting date and time (FROMTIME)

To actually copy a file, it must meet all of the above three requirements.

You can use CPYOPT to perform a backup of your optical volumes, but is not the suggested way of doing so. Remember that CPYOPT works on a file basis, therefore if you are copying a large number of files your CPYOPT request could take hours to complete. What options you specify can also affect how long your copy request could run. Review the examples later in this chapter for a comparison of your options and how they may affect your copy request.

Select files to copy

Use the **Select files to copy** (SLTFILE) parameter to choose which files you want to copy. Your options are:

- ***CHANGED** is the default option. The system will determine if the source file already exists on the target volume. If so, the system will copy the source file only if the source file has changed since performing the last copy. Two sets of dates and times determine if a file has changed: Either the date and time the file was last changed, or the date and time the file attributes were last changed. Specifying DATA(*FILATR) on the Display Optical (DSPOPT) CL command can display these dates and times.
- ***ALL** indicates that the system will copy all files on the source volume. The system will replace any files which may already exist on the target media by a new copy from the source volume.
- ***NEW** indicates that the system will copy only files which are not currently on the target volume.

If the target volume already contains files, choosing the ***CHANGED** or ***NEW** option may result in a longer running CPYOPT request. This is because the system has to make a list of files for both the source and target volume, and then compare them. The time required to do this can become excessive when the volumes contain thousands of files.

Copy subdirectories

Use the **Copy subdirectories** (CPYSUBDIR) parameter to indicate whether or not to process files in the subdirectories of the specified **From path**. Your options are:

- ***NO** indicates that only files in the specified **From path** are eligible to be copied. This is the default option.
- ***YES** indicates that files in all subdirectories of the specified **From path** are eligible to be copied. The system creates subdirectories on the target volume if they do not already exist. The newly created subdirectories have the same name as they did on the source volume, even though the parent directory name can be

different. A system makes a check prior to the copy operation to ensure that any resulting new path name does not exceed the maximum path name length. The system prevents you from copying the subdirectories of one directory to a subdirectory of that directory on the same volume.

Starting date and time

The system will use the FROMTIME parameter to determine if a file is eligible for copying based on its creation or modification date. All files that were created, changed, or whose attributes have changed, on or after the starting date and time are eligible for copying. You can determine when a file was last created or changed by specifying DATA(*FILATR) on the Display Optical (DSPOPT) CL command. The default values: *BEGIN for **Starting date** and *AVAIL for **Starting time** indicates that all files meet the starting date and time requirement. Specifying a starting date and time identifies only files that were created or changed since that date and time as eligible for copying. **You can use this parameter to greatly limit the number of files that require processing by CPYOPT. This decreases the time that is required to process the files.** You could combine this parameter and the SLTFILE parameter to limit the number of files that need to be checked before copying. You could select only files that were *CHANGED or *NEW after a specified starting date and time.

Examples

Scenario 1 — Copy all files from the source volume: This example shows how to copy all files from the source volume VOL001 to a volume which currently does not contain any files or directories. The system will process all subdirectories of the source volume, create the subdirectories on the target volume, and copy all files.

```
> CPYOPT FROMVOL(VOL001) FROMPATH(/) TOVOL(CPYVOL001) SLTFILE(*ALL)
   CPYSUBDIR(*YES) CRTDIR(*YES)
```

Scenario 2 — Copy all files from the source volume since the last copy request:

For this example you have several different option which may take different lengths of time. Your first option would be to issue the same request as Scenario 1 but with a different target volume. The system will copy all the files and directories to the new target volume.

Your second option would be to use the *CHANGED option on the SLTFILE parameter.

```
> CPYOPT FROMVOL(VOL001) FROMPATH(/) TOVOL(CPYVOL001) SLTFILE(*CHANGED)
   CPYSUBDIR(*YES) CRTDIR(*YES)
```

Depending on how many files are currently on the source and target media this request may take a long time to process. First, you must obtain a list of all files on the source media and the target media. Then compare the files in order to determine if any file has changed since the last CPYOPT request. Once that is done, the system will copy only the files that have changed.

The third option is using the *NEW option on the SLTFILE parameter, provided that no existing files changed, but simply added to the source volume.

```
> CPYOPT FROMVOL(VOL001) FROMPATH(/) TOVOL(CPYVOL001) SLTFILE(*NEW)
   CPYSUBDIR(*YES) CRTDIR(*YES)
```

First, you must build option a list of all files that are on both the source and the target volume. Then compare the files before copying any new files.

Your fourth option would be to use one of the SLTFILE options in combination with specifying a starting date and time.

```
> CPYOPT FROMVOL(VOL001) FROMPATH(/) TOVOL(CPYVOL001) SLTFILE(*CHANGED)
    CPYSUBDIR(*YES) CRTDIR(*YES) FROMTIME('04/01/99' '00:00:00')
```

By specifying a starting time, the system copies only files that were created or changed after 04/01/99 to the target volume.

Using CPYOPT to Backup a volume

You can use the Copy Optical (CPYOPT) command to do a complete or partial copy of your volume. The following describes the special processing of the CPYOPT command when the target is a *PRIMARY or *BACKUP volume. Refer to the following list to decide how to best use the command.

- Copy completely or partially. You can copy a file, a directory, a directory with all of its sub directories, or an entire volume.
- Copy incrementally. You can copy only what has changed since the previous CPYOPT request.
- Copy specifying a starting date for selecting files. Only files that are created or changed on or after the specified date are eligible for copying.
- Replicate the hierarchical structure of the source volume on the target volume.

CPYOPT requirements to volume type *PRIMARY

When the target volume is type *PRIMARY, the CPYOPT command has the following unique requirements:

- The source volume can be either type *PRIMARY or *BACKUP.
- Because the target volume is *PRIMARY, all API requests and most optical commands and utilities can access the volume.
- Because utilities and user programs can update the *PRIMARY volume, you need to determine how to protect directories and files from unauthorized change or deletion.
- The target volume could contain information for one or multiple optical *PRIMARY volumes. An easy way to manage multiple volumes on a single target volume would be to have a new first level directory. That directory name could be the name of the source-primary volume.
- You need a way of keeping track of when a volume or directory was last backed up. Use the CPYOPT command to do it automatically.
- The hierarchical structure on the target volume does not need to be identical to that of the optical *PRIMARY volume.
- The create date and time, and change date and time, of the file on the target volume will be different than their counterparts on the optical primary volume. The file creation data and time on the target volume is the date that the file was written.
- You can use directories and files on the target *PRIMARY volume directly. You do not have to copy applications back to a *PRIMARY optical volume.
- You can request that the system copy only new files on the source volume to the target volume. This might be useful if you never changes files on your source volume but only create new ones.

CPYOPT requirements to volume type *BACKUP

When the target volume is type *BACKUP, the Copy Optical (CPYOPT) command has the following unique requirements:

- The source volume must be type *PRIMARY.

- Only the CPYOPT and CVTOPTBKU commands can write to the target *BACKUP volume. APIs, utilities, and user programs cannot write to an optical *BACKUP volume.
- An optical *BACKUP volume can contain information for only one optical *PRIMARY volume. This prevents two primary volumes from sharing the same optical backup volume.
- You cannot delete directories or files from an optical *BACKUP volume. This ensures data integrity of the optical backup volume.
- The system maintains the file create date and time, and change date and time for the *PRIMARY volume on the optical *BACKUP volume.
- A user application cannot directly use a file or directory on a *BACKUP volume. First, you must copy the file or directory to an optical *PRIMARY volume by using CPYOPT.
- If the optical *PRIMARY volume is damaged or lost, you can convert the optical *BACKUP volume to an optical *PRIMARY volume. You can do this by using the Convert Optical Backup (CVTOPTBKU) CL command.
- To maintain control information about the status of backup requests, optical *BACKUP volumes require additional media usage. Because of this, a *PRIMARY volume that is 100% used may not fit on an optical *BACKUP volume.
- The system always initializes *BACKUP volumes with a 99% volume threshold.

Suggestions on which CPYOPT volume type to use

Here is a list of items that you can use to determine if you should use a target volume type of *PRIMARY or *BACKUP.

- In general, the CPYOPT to a type *PRIMARY volume gives you more flexibility, but it requires more management of your backup volumes.
- The CPYOPT to a type *BACKUP volume provides more management and security for your optical backup volumes, but it is less flexible.
- Use the CPYOPT command, and specify a type *PRIMARY volume, if you want to copy data from several volumes to a single volume.
- Use the CPYOPT command, and specify a type *BACKUP volume, if you want better security for your backup volumes. The system cannot write volumes with type *BACKUP to with normal optical commands or user programs.
- Use the CPYOPT command and specify a type *BACKUP volume to save information such as when the system copied directories and volumes, and the success status of those copies.
- The biggest advantage of using the CPYOPT to a type *BACKUP volume is that the system stores the backup control information on the backup volume. This information includes the relationship between files on the backup volume and the files on the primary volume. This information can be very useful if you ever need to recover the lost source data is from the backup volume.
- Use the CPYOPT command, and specify a type *BACKUP volume, if you want the source and target file dates (creation and change) to be identical.
- One disadvantage in using the CPYOPT command to a type *BACKUP volume is that the system uses extra space on the backup volume to store control information. The amount that is used is approximately 3 times the sector size for each directory. Therefore, if CPYOPT copies 100 directories from a primary volume to a backup volume where the sector size on the backup volume is 2KB, the backup volume uses an additional 600KB of space. In this example 600KB of space is used every time the command is run.

Volume type *BACKUP

This section discusses the *BACKUP volume type and the unique backup process that are associated with a *BACKUP volume. Remember that using Copy Optical (CPYOPT) to a *BACKUP volume is no longer the suggested way of backing up your volumes because of the time required. Duplicate Optical (DUPOPT) is the recommended way of backing up your volumes.

The *BACKUP volume type supports the backing up and protection of information from optical *PRIMARY volumes. The system does not allow user programs or APIs to write to optical *BACKUP volumes. Only a limited set of optical commands can update *BACKUP volumes. Once the system creates a directory or file on an optical *BACKUP volume, the only way to delete them is to reinitialize the volume. Doing this prevents either accidental or intentional deletion.

Backup volumes and directories contain a **Complete Backup Range** file which contains date information about prior copy requests to the optical *BACKUP volume. These dates are helpful in determining the contents of backup directories and volumes with respect to the contents of their primary counterparts. These control dates make it easier to recover by providing a time checkpoint. Each backup directory has its own control dates. Each backup volume also has its own control dates, which include:

- Complete starting date and time
- Complete ending date and time
- Last changed date and time

The system writes these dates to the backup volumes in a reserved file within each backup directory. Since the system writes the dates to the media, the backup volumes are self-contained. Not only is the backup data on the media, but the recovery information is there as well.

Complete Backup Range

What is a **Complete Backup Range**? When an optical *PRIMARY volume is copied to an optical *BACKUP volume a special file called a **Complete Backup Range** is written to the *BACKUP volume. This file indicates the last time a backup was done. The system keeps backup control information for the volume as well as each directory on the volume. If the volume or directory was successfully backed up the **Complete Backup Range** will contain both a starting and ending date and time. When a range exists for an optical backup volume or directory, it has a specific meaning: The backup directory or volume has a copy of all the created or changed files within the date range that correspond to the primary directory or volume.

For example, volume BVOL1 is an optical *BACKUP volume for *PRIMARY volume PVOL1. BVOL1 contains directory /DIR1 that has a **Complete Backup Range** as follows:

```
Start Date: 1/1/99
Start Time: 09:00:00
End Date:   1/30/99
End Time:   22:00:00
```

This means that the system backed up all the changed or created files in /DIR1 on PVOL1 since 9:00 a.m. on 1 January 1999. The system backed up files to /DIR1 on BVOL1 at 10:00 p.m. on 30 January 1999. Any files that were created or changed on *PRIMARY volume PVOL1 in directory /DIR1 after 22:00:00 on 1/30/99 would not yet have been backed up.

Complete Backup Range for Directories: The **Complete Backup Range** for a directory does not encompass all subdirectories of the directory. In other words, each directory has its own unique **Complete Backup Range**. For example, assume that directory /A has a **Complete Backup Range** of 1 March 1999 through 1 May 1999. This does not necessarily mean that directory /A/B has the same complete range. In fact, /A/B might have no **Complete Backup Range** at all. The complete range does not reflect a hierarchical range over all directories within that subtree.

The system updates the **Complete Backup Range** for a backup directory after it copies all eligible files in the primary directory.

Note: Use the SLTFILE parameter on the CPYOPT command to determine if a file is eligible. If you use *ALL, all files are eligible to copy. If you use *CHANGED, only those files that were created or changed since the last CPYOPT command are eligible. If you specified *NEW, the system copies files only if they do not exist on the target volume.

For example, FILE.001 gets copied on 1 March 1999 as a result of a complete backup of directory /DIR1. At this time /DIR1 is given an ending range of 1 March 1999. On 1 April 1999, the user has the system back up directory /DIR1 again by specifying SLTFILE(*CHANGED). However, the back up affects only the files that have changed. If FILE.001 has not changed since the previous CPYOPT command, this file is not eligible to copy. However, the system updates the ending range for /DIR1 to 1 April 1999 if none of the eligible files fail to copy.

Complete Backup Range for Optical Volumes: The **Complete Backup Range** for an optical volume is very similar to that of an optical directory. The complete range for a directory represents the relationship between the files in a backup directory and those in the primary directory. Likewise, the complete range for an optical backup volume represents the relationship between the files on an optical backup volume and those on the primary volume. You must back up all eligible files on a volume to update the complete range of the volume.

You can update the complete range for a volume only if the CPYOPT command specifies the FROMPATH(/) and CPYSUBDIR(*YES) variables. This ensures that the system will process all files on the *PRIMARY volume.

Complete Backup Range — Starting Date and Time

The **starting date and time** of a **Complete Backup Range** for an optical backup volume or directory is the earliest time that is specified on a CPYOPT command when all eligible files on the volume or directory were successfully copied. You can specify a starting date and time on the CPYOPT command. This system uses time to select the files from the primary volume to copy to the optical backup volume. The system copies any files that are created or changed on or after this time. First, the system must successfully copy all eligible files for a directory or volume. Then the system sets the starting date and time for the corresponding optical backup volume or directory to the specified time. The definition indicates that this value is the earliest time that is specified on a CPYOPT command. Consider the following example.

Starting Date and Time — Scenario: A user issues the CPYOPT command for directory /DIR1 by specifying 1 May 1999 as the starting date. If all eligible files successfully copy, then the system sets the complete starting date for the backup directory /DIR1 to 1 May 1999.

Now assume that the user issues the CPYOPT command again for /DIR1. This time the system sets the starting date to 1 April 1999. This request copies any files that have changed since the last CPYOPT command. Additionally, it copies any files that were created between 1 April 1999 and 1 May 1999 that were not selected on the previous request. If all eligible files again copy successfully, then the starting date for backup directory /DIR1 changes to 1 April 1999. Future copies specifying earlier starting dates would produce similar results.

Use *BEGIN and *AVAIL for the starting date and time on the CPYOPT command. This will copy all the files from a primary directory or volume, regardless of the create or change time for the file.

Complete Backup Range — Ending Date and Time

The CPYOPT command does not allow you to specify an ending date and time. The system always uses the date and time of the copy request as the ending date and time. Therefore, the system uses the date and time of the request for the complete ending date and time for a backup directory or volume.

The ending date and time of a **Complete Backup Range** for an optical backup volume or directory is one of the following conditions:

- The last time a CPYOPT command completed.
- When the system successfully copies all eligible files in that volume or directory.
- When the starting date and time of the request is not after the existing complete range.

The definition of the ending date and time field has two parts. First, this date is the last time a CPYOPT command completed for the directory or volume with no failures. Second, the complete ending date and time will not update if the range of the request does not overlap the existing range. This is true even if all eligible files copy successfully.

Ending Date and Time — Scenario: On 1 July 1999, the user issued the CPYOPT command for directory /DIR1 that specifies 1 February 1999 as the starting date. If all eligible files successfully copy, the system sets the complete starting date for the backup directory /DIR1 to 1 February 1999. The system sets the complete ending date to 1 July 1999.

Now, the system issues a second CPYOPT command for directory /DIR1 on 15 September 1999, specifying 1 June 1999 as the starting date. If all eligible files successfully copy, the complete starting date for backup directory /DIR1 remains 1 February 1999. The complete ending date is moved out to 15 September 1999. This is the normal situation that takes into account only the first part of the definition above.

On 1 December 1999, the user issues the CPYOPT command again for the /DIR1 directory. This time the user specifies 1 October 1999 as the starting date. Even if all eligible files copy successfully, the complete range will not change. The complete range cannot be expanded to include the new ending date. This is because the files that were created or changed between 15 September 1999 and 1 October 1999 are not accounted for.

Complete Backup Range — Lasted Changed Date and Time

Using the CPYOPT command causes the system to write the **last changed date and time** of an optical backup volume or directory.

This includes any time that the system wrote files or directory attributes to the directory or volume.

The last changed date and time for that directory and volume will always reflect the date and time of the request. This remains true even if the system writes a file to a backup directory.

Last Changed Date and Time — Scenario One: On 1 July 1999, the user issues the CPYOPT command for directory /DIR1 by specifying *BEGIN as the starting date. If the system successfully copies all the eligible files, then the dates are as follows:

- The system sets the complete starting date for backup directory /DIR1 to *BEGIN.
- The system sets the complete ending date to 1 July 1999.

If the system copied at least one file to /DIR1 as a result of this request, the last changed date is also 1 July 1999.

The system does not necessarily update the last changed date and time as the result of a successful copy. If the system did not write any files to the backup directory, the system may update the complete range, but not the last changed date.

Last Changed Date and Time — Scenario Two: In “Last Changed Date and Time — Scenario One,” the backup directory /DIR1 has the following as dates after the request:

- A starting date of *BEGIN.
- An ending date of 1 July 1999.
- The last changed date of 1 July 1999.

On 1 October 1999, the user issues the CPYOPT command again for directory /DIR1. This time the command specifies SLTFILE(*CHANGED) to copy only the files that have changed since the last CPYOPT request. Assume that no files have changed since the last backup on 1 July 1999. Since no files are eligible to copy, the system writes no files to the backup directory /DIR1. Therefore, the last changed date remains 1 July 1999. However, since no eligible files failed, the complete range for /DIR1 expands to have an ending date of 1 October 1999.

The last changed date and time becomes most important when it is set beyond the complete range. This would happen if some files were actually copied but other eligible files failed to copy for some reason.

Last Changed Date and Time — Scenario Three: In “Last Changed Date and Time — Scenario Two,” the backup directory /DIR1 has the following dates after the request:

- A starting date of *BEGIN.
- An ending date of 1 October 1999.
- A last changed date of 1 July 1999.

On 1 December 1999, the user issues the CPYOPT command again for directory /DIR1. Assume that 10 files were changed or added to primary directory /DIR1 since the last CPYOPT request on 1 October 1999. Assume that only eight files successfully copy to /DIR1 and that two of the eligible files failed. Since the system did not copy all eligible files, the complete range stays the same with a starting date of *BEGIN and an ending date of 10/1/99. However, since /DIR1 changed, the last changed date gets updated to 1 December 1999. Since the last

changed date is outside the complete range, a complete copy of /DIR1 from *BEGIN to 1 October 1999 may not exist. A more recent copy by the change on 1 December 1999 might have replaced one of those files.

Copying to Optical *BACKUP Volume — Example

This example shows the system that backs up the *PRIMARY volume VOL01 to the *BACKUP volume BKP-VOL01. This will copy all the files in all the subdirectories. Once the system writes to volume BKP-VOL01, the system will use the volume in one of the following ways:

- For further backups of volume VOL01.
- For converting from *BACKUP volume BKP-VOL to *PRIMARY volume VOL01.
 - > CPYOPT FROMVOL(VOL01) FROMPATH(/) TOVOL('BKP-VOL01' *BACKUP)
SLTFILE(*ALL) CPYSUBDIR(*YES)

Converting an Optical *BACKUP Volume

Use the Convert Optical Backup (CVTOPTBKU) command to convert an optical *BACKUP volume to an optical *PRIMARY volume. You would typically use this function when the *PRIMARY optical volume is either damaged or missing. The conversion eliminates the necessity of copying all information from the optical *BACKUP volume to a new *PRIMARY volume. After the system converts the volume to a *PRIMARY volume, it will allow all write requests to the volume.

Once the system converts an optical *BACKUP volume to a *PRIMARY volume, there is no way to convert it back to an optical *BACKUP volume. To convert backup volumes, select option 6 (Convert optical backup volume) from the Optical Backup/Recovery display or use the CVTOPTBKU CL command.

Before you attempt to convert, you should verify the name of the *PRIMARY volume for which this volume is a backup. You can do this by displaying the volume attributes of the optical *BACKUP volume. You can do this by using the Display Optical Volume Attributes (DSPOPT) command or by selecting option 5 (Display) from the Work with Volumes display.

There may be previously deleted *PRIMARY volume directories and files on the optical *BACKUP volume. Therefore, when converting the optical *BACKUP volume to a *PRIMARY volume, it might be necessary to manually delete directories and files from the volume. You perform this function in order to accurately reflect what was on the *PRIMARY volume. If you never delete directories and files from a *PRIMARY volume, this should not be a concern.

CPYOPT Performance

Performance is a complex subject with many interdependent components. By changing one component, you may adversely affect another. This interdependence and the other factors that affect performance, prohibits providing a formula for computing the time required to copy a given number of files. It is simply a guideline to help you estimate how long your CPYOPT might take.

You may need to estimate how long it takes to copy an entire optical volume when using the CPYOPT command. You can start by copying a single directory that contains a known number of average sized files. Then take the difference between the ending time and starting time to determine an elapsed time. Take the elapsed time and divide by the number of files that are copied to figure the average seconds per file. You can use this number as a basis to determine the amount of time that is required to copy the entire volume of average size files.

To maximize copy performance, use the following set of guidelines as a starting point:

- Having too few directories with too many files can affect performance. Having too many directories with too few files can affect performance also. Try to keep the number of files in a directory to less than 6000 files.
- Consider performance when determining file size.
- Avoid use of extended attributes on files. When a file has extended attributes, they are stored separately from the data. When copying the data, they must be copied also. It is similar to copying a second file for each user file copied.
- Keep the source and target volumes in the same library.
- Avoid copying to the opposite side of an optical cartridge.
- If the copy processes can have dedicated use of the optical drives, use the COPYTYPE *IOP parameter on the CPYOPT command.
- Avoid optical drive contention from other optical processes.
- Dedicate the use of two optical drives for copy activity.

Drive Contention

The following conditions can severely affect copy performance:

- Having only one drive available for use.
- Copying from one side of an optical cartridge to the opposite side.
- Having other optical processes that are running that attempt to use the available drives.
- A large number of files on source volume.

To remove a volume, to store it in a slot, to retrieve a new volume, and to mount it requires from 8-15 seconds. You should try to do your copy requests when the process can have dedicated use of the optical drives.

You should not try to copy a large number of files from one side of an optical cartridge to the other side. Optical drives have only one read/write head. The following conditions occur when copying from one side of an optical cartridge to another:

- The system mounts the source volume.
- A limited number of files that are to be copied are read and stored on OS/400 temporary storage.
- The source volume is removed, and the system mounts the target volume by turning over the optical cartridge.
- Files are read from OS/400 temporary storage and written to the target volume.
- If there are more files to copy, the system removes target volume and mounts the source volume again by turning over the optical cartridge.
- The system repeats this process until it copies all the files. You may need to turn the optical cartridge over many times to copy all the files.

Copy performance is always better when both the source and target volumes are in the same library and when setting the COPYTYPE parameter for that library controller to *IOP. Two conditions exist that require an extra processing step. The first one is that the source and target volumes are in different libraries. The second involves setting the COPYTYPE parameter to *SYSTEM and having the volumes exist in the same library. This extra step requires moving the files you want to copy to temporary storage on the iSeries server before writing them to the target volume. You do not need to use temporary storage when the system meets both of the following conditions:

- Both optical volumes are in the same library.
- You set the COPYTYPE parameter on the CPYOPT command to *IOP.

This allows for the direct transfer of data between the two optical drives.

SAV/RST

The SAV command can be used to create a backup of an optical volume image. An optical volume image is a copy of the entire optical volume in *SAVRST format. Using SAV, the volume image can be saved to any supported save/restore device including tape, diskette, optical, or save file. Subsequently when the volume image is restored using the RST command, the entire image must be restored to an existing optical volume either in a stand-alone device or an optical media library.

An optical volume image has unique properties that require the entire volume image to be saved or restored in a single operation. Once saved you cannot restore individual files or directories.

Once an optical volume image is saved, it can be viewed with DSPTAP, DSPDKT, DSPOPT, or DSPSAVE, depending on the save/restore device used. When the volume save/restore entry is displayed, option 8 can be used to display the additional information panel, which includes media specific information such as media type, volume capacity, sector size, and security attribute information. You cannot see the individual files and directories that make up the volume image.

Use of the generic SAV command to save optical data can be easily incorporated into an existing system backup strategy without requiring a separate command such as DUPOPT to perform the save operation. SAV provides a good alternative to DUPOPT because it allows a volume to be saved from a one drive optical media library or from a standalone device without requiring the allocation of a second optical device. SAV provides a viable incremental backup solution by periodically backing up volumes not yet at capacity to a save/restore device such as tape. When the volume is full it can be duplicated for archival purposes by either restoring the full volume to create a copy or by issuing DUPOPT to duplicate the volume.

Saving and then restoring an optical volume image creates an exact copy of the saved volume including the volume name. DUPOPT creates a copy of the source volume but the volume name is changed.

To save and restore an optical volume the following authority is required:

- *USE authority to the optical device.
- *SAVSYS special authority OR *OBJEXIST authority through the optical volumes authorization list.
- If the media format is UDF *RWX authority is also required to the root directory of the volume. The device will be locked shared (LSRD) read while a SAV or RST is active.

Auditing records created during a SAV or RST request.

- OR** Object Restored
- RZ** Change primary group during restore. Saved value different from target. Value on target remains unchanged. (UDF only)
- RO** Change owner during restore. Saved value different from target. Value on target media remains unchanged. (UDF only).

- O1 Successful open for save (S/R/S) Storage/Read/Save. Successful open for restore (S/U/R) Storage/Update/Restore

SAV Support

SAV can be used to save HPOFS or UDF (Universal Disk Format) formatted volumes. This function will not support the backup of a ISO9660 formatted media.

Select a volume or volumes to save. The SAV command prevents the implicit saving of all optical volume images in the QOPT file system when the OBJ parameter includes the entry '/*'. The file system QOPT cannot be saved; however, volumes below the file system can be saved. If you want to save all volumes within the QOPT file system, '/QOPT/*' must be explicitly specified on the OBJ parameter. If all volumes are selected, be aware that this SAV operation could take a long time to complete depending on the number of volumes being saved. See the "Parameters" section for additional restrictions on the OBJ parameter.

In order to specify that a volume image is to be saved, you must specify a value of *STG on the SUBTREE parameter.

Saving an optical volume image to another optical volume is allowed; however, the target volume cannot be the opposite side of the saved volume.

The performance of SAV is comparable to DUPOPT, although it depends on the target device chosen.

Parameters

OBJ Specify a single or multiple path names. The path name cannot be extended beyond the volume level. Examples of invalid path names include, '/QOPT/VOL/*' or '/QOPT/VOL/DIR/FILE'.

SUBTREE

Must be *STG when saving optical volume images.

CHGPRIOD

Start Date, Start time, End Date and End time must all be *ALL.

UPDHST

Must be *NO.

SAVACT

Parameter is ignored when attempting to save optical volume images.

PRECHK

Must be *NO.

TGTRLS

Value cannot precede V5R2M0.

Refer to the CL topic in the Programming category in the Information Center for a further explanation of parameter values and their meaning.

Examples

- Save all volumes within the QOPT file system to a save file.
SAV DEV('/qsys.lib/xyz.lib/xzysavfile.file') OBJ('/qopt/*') SUBTREE(*STG)
- Save all volumes beginning with vola and volb to a save file.
SAV DEV('/qsys.lib/xyz.lib/xzysavfile.file') OBJ('/qopt/vola*') ('/qopt/volb*') SUBTREE(*STG)
- Save one volume vol1 to a tape device.

```
SAV DEV('/qsys.lib/tap01.devd') OBJ('/qopt/vol1') SUBTREE(*STG)
```

RST Support

RST can be used to select a volume or volumes to restore. The RST command protocol requires that all physical file systems adhere to a certain set of predefined rules governing how restore is handled, depending on whether or not the object exists on the system. For purposes of restoring an optical volume image, the target media must exist on the system, either mounted in a stand-alone device or imported into an optical media library. Also, it must be accessible by the name specified on the **OBJ-New Path Name** parameter. The **OBJ-New Path Name** must either match the name of the OBJ-Name parameter or be *SAME. This restriction will require that unformatted volumes be renamed before executing the RST command. Unformatted volumes can be renamed using option 7 from either the WRKOPTOL, WRKLNK display or by issuing the generic RNM command.

When an unformatted volume is renamed, the new name acts as an alias to the unformatted volume. The new name will not be written to the media and will not be preserved if the volume is removed from the device. The name is only a temporary volume identifier used to refer to this volume until the volume is restored.

Saved volumes can be restored to both unformatted and formatted volumes. If restoring to a formatted volume that contains active files and directories, an inquiry message is sent. If you proceed with the restore, all data on the target media will be lost.

Saved HPOFS volumes can be restored to Erasable media with matching sector sizes and a capacity equal to or greater than the saved volume.

Saved UDF volumes on DVD and Erasable can be restored onto DVD or erasable media, but media capacity and sector size must be identical to the saved volume.

WORM volumes can be restored to either WORM or Erasable media as long as the capacity of the target media is greater than or equal to the saved volume capacity and the sector size is identical to the saved volume. When restoring to WORM, the target volume must be unformatted.

The performance of RST is comparable to DUPOPT, although it depends on the target device chosen.

Parameters

OBJ - Name

Name of the optical volume image or images to be restored from a save/restore device.

OBJ - New Path Name

Specify a single or multiple path names. The path name cannot be extended beyond the volume level. Examples of invalid path names include, '/QOPT/VOL/*' or '/QOPT/VOL/DIR/FILE'. Specify the names of existing volumes or *SAME.

SUBTREE

Must be *STG when restoring optical volume images.

OPTION

Must be *ALL or *OLD.

ALOWOBJDIF

Select *OWNER, *PGP, *AUTL, *NONE, *ALL. The value selected determines what differences will be tolerated between the saved volume and the formatted target volume. If changes are allowed, an attempt will be made to preserve the security attributes UID, GID and PERMS of a the UDF target volume root directory. The *OWNER and *PGP values are not checked when restoring to an uninitialized volume or when restoring to an initialized HPOFS volume.

Refer to the CL Reference guide for a further explanation of parameter values and their meaning.

Examples

- Restore all volumes within the QOPT file system from a save file.
RST DEV('/qsys.lib/xzylib.lib/xzysavefile.file') OBJ((* *INCLUDE *SAME)) SUBTREE(*STG)
- Restore all volumes beginning with vola and volb from a save file.
RST DEV('/qsys.lib/xzylib.lib/xzysavefile.file') OBJ('/qopt/vola*' *INCLUDE *same) ('/qopt/volb*' *INCLUDE *same) SUBTREE(*STG)
- Restore one volume, vol1 to vol1.
RST DEV('/qsys.lib/tap01.devd') OBJ('/qopt/vol1' *INCLUDE *same) SUBTREE(*STG)

Note: The OBJ-New Path Name must either match the name of the OBJ-Name parameter or be *SAME. This restriction will require that unformatted volumes be renamed prior to executing the RST command.

Chapter 9. Optical Media Library Performance

This chapter describes performance considerations for users of optical media libraries (both directly attached and LAN-attached).

General Performance Considerations for Optical Support

Several factors can affect the optical performance of both LAN-attached and directly-attached optical media libraries. This topic explains how these factors can influence optical performance.

To view performance data for optical devices, go to the "Performance" link in the iSeries Optical Storage home page.

Volume Mounting and Dismounting

Volume mounting and dismounting is one of the most important factors affecting optical performance. It takes approximately 8 to 15 seconds to remove a volume, store it in a slot, retrieve a new volume, and mount it. If you can minimize the number of volume mounts and dismounts that your application requires, optical performance will improve.

Drive Contention

Performance can be severely affected by drive contention. The following conditions increase drive contention and should be avoided:

- Having only one drive available for use by applications libraries.
- Having many optical processes running that attempt to use different optical volumes at the same time.

Number of Directories and Files

Performance can be affected by having too few directories with too many files. Directories group related information to provide a means of quicker access. Typically, you get better performance from more directories with fewer files. Although there is no enforced limit on how many files there can be in a directory, you probably should not have more than 6000 for performance reasons.

File Size

The size of a file has a direct effect on the amount of time it takes to read, write, or copy the file. In general, the larger the file, the longer the operation can be expected to take. For more information on the effects file size have on performance, see "Working with Files" on page 87.

Performance Considerations—Directly-Attached Optical Libraries

This topic addresses performance considerations that are specific to directly-attached optical libraries.

Effects of File Attributes on Performance

When a file has extended attributes, they are stored separately from the data. When the data is written or copied, the attributes must also be written or copied. If file attributes are not required, attribute copying can be suppressed when copying

between the QOPT and QDLS file systems by using the Change Optical (CHGOPTA) command. Setting the copy attributes (CPYATR) value on the CHGOPTA command to *NO suppresses the copying of attributes between the QOPT and QDLS file systems.

HFS API Expanding Buffer I/O

Users of the HFS APIs can improve performance by taking advantage of the expanding buffer I/O option. Expanding buffer I/O lets you control the amount of data that is read from the optical media when only parts of the entire file need to be read. For information on using expanding buffer I/O, see “Expanding Buffer I/O through HFS—Tailoring Read Requests for Performance” on page 84.

CPYOPT and DUPOPT Performance Considerations

Performance of the CPYOPT and DUPOPT functions is a complex subject with many interdependent components. By changing one component, you may adversely affect another. Because of this interdependence and the other factors that affect copy and backup performance, you should read Chapter 8, “Optical Volume Backup,” on page 89 to better understand how to achieve the optimum performance using the CPYOPT and DUPOPT commands.

Volume Mounting and Dismounting Schedule Timers

This topic provides information on how the iSeries server manages work requests to the directly-attached optical library devices. It is a high-level view and does not include all program logic.

You can use the Change Device Description (CHGDEVMLB) command to change the queuing and scheduling logic used by the iSeries server for directly-attached optical media libraries.

There are two timer values that are associated with optical media libraries that affect the scheduling of volume mounting and pre-emptive dismounts. You can change both timer values by using the CHGDEVMLB command. The first timer value (UNLOADWAIT) is the unload wait time. This value determines how long the system waits for a new request for a mounted volume before removing it. The second timer value (MAXDEVTIME) is the maximum device wait time. This value determines how long a volume with active requests remains in a drive while other queued requests are waiting to use the drive.

By using these two timer values, you can adjust the volume mount scheduling that is used by the optical media library to match your application’s use of optical volumes.

You can change these timer values at any time; however, the new timer values will not become effective until the next time the device is varied on.

System job priority and limit timers are used to schedule volume mounting. The maximum number of volumes that can be mounted is equal to the number of drives in the optical media library. Keep the following points in mind as you schedule volume mounting:

- A volume can remain mounted in an optical drive for the maximum device wait time if work requests with the same or lower job priority for a different volume have been received. An exception to this is when you are initializing a rewritable volume or using DUPOPT; the volume remains in the drive until the operation completes.

- Work requests for mounted volumes are serviced before requests of the same or lower job priority for volumes not mounted.
- For a multiple job environment, volumes are mounted based on the job priority for the work request. A work request from a job with a higher system job priority causes the required volume to be mounted to handle that request. The volume remains mounted for the maximum device wait time if work requests continue, dismounts after unload wait time inactivity, or is pre-empted by a work request from a job with higher system priority. If you are initializing a rewritable volume or using DUPOPT, the volume remains mounted until the operation completes.
- If the work on a drive is interrupted because of a higher priority request, the maximum device wait time timer for the currently mounted volume is cancelled. All future requests for that volume are queued for normal processing by priority.
- If the volume needed for a work request is not mounted within the system timeout (typically 30 minutes), the job fails due to a time-out.

Changing Job Priority on Active Jobs

Due to the work management method used by OS/400 optical support, changing run priorities of an active optical job at the user level can result in loss of time allocation and, in some cases, cause jobs to time out.

Add Optical Cartridge Performance

Specifying *NO for Rebuild Directory Index can improve the performance of ADDOPTCTG by deferring the build of the optical directory index until a later time.

Performance Considerations—LAN-Attached Optical Libraries

This topic describes the ways you can control the performance of applications using an optical LAN server.

Controlling the Block Size for an Optical LAN Server

Optical LAN support is optimized to receive and send the largest possible block of data across the network. Currently, optical LAN support communicates with servers using data frames containing a maximum of 31KB bytes of information. This does not directly affect the HFS user, but to minimize communications time, applications using large files should avoid read and write operations of less than 31KB. Read or write operations larger than 31KB are not a problem because they are automatically broken up into the optimum size.

Controlling Block Size for Optical LAN Server—Examples

Consider an example application that has 64KB bytes of information to be written to an optical volume. The application could be designed to write blocks of 1KB, 31KB, or the entire file. If 1KB blocks were written, the application would make 64 write requests. These 64 write requests would result in 64 communications requests.

If 31KB blocks were written, the application would make 3 write requests (2(31KB) + 1(2KB)). These 3 write requests would result in 3 communications requests.

Finally, if the application requested that the entire file be written, there would be 1 write request. The file would automatically be broken up into 3 communications requests (2(31KB) + 1(2KB)).

All three approaches work, but the best performing one is the single write request of the entire file.

While a single write request may be appropriate for some applications a single read request may not. To correctly optimize a read application, you first need to determine what portion of the file is required. Assume that the the same 64KB file in the example above is a read application, but that you will most likely need only 1KB of information. If the application requests that the entire file be read, this would result in three communication requests (2(31KB) + 1(2KB)). If the application requested a 31KB read request, this would result in a single communication request. A 1K application read request would also result in a single communications request. Even though the 31K and 1K read request have the same results, you might be better off with the 31K read request if you anticipate more information will be needed by the application.

Conversation Allocation for Optical LAN Server

You have performance control over how files are opened and closed, and how requests to optical volumes are issued from the application. Either of these can result in excessive communications overhead.

But, before addressing that, it is important to understand how conversations are allocated and de-allocated between OS/400 and the optical LAN server. OS/400 and the optical LAN server can support a limited number of conversations between them. The mode status of the iSeries server determines how many conversations can be started and owned by the system.

The optical LAN server has its own mode status that determines how many conversations it can own and have active. A limit is negotiated between the system and the optical LAN server when communications are initially started. To view what that limit is, use the Display Mode Status (DSPMODSTS) command for the device that defines the optical LAN server. From the displayed list of modes, display the details for the mode that is used by the device. Locate the negotiated session limits. This number is the maximum number of concurrent applications that can use the LAN server. Different servers can have different limits.

This negotiated limit can be viewed as a pool of conversations. To stay within this conversation pool limit, optical LAN support was designed to use and return conversations as needed. It is also designed to keep a conversation as long as the application has at least one open file.

Conversation Allocation for Optical LAN Server—Example

Following is an example application and information on how you could change it to optimize conversation allocation. This sample application performs the following actions:

1. Find the available space on a volume
2. Create a directory
3. Open file 1
4. Write to file 1
5. Close file 1
6. Open file 2
7. Read from file 2
8. Close file 2

Running the example application would result in 4 different conversations. The first conversation is allocated for the volume space request. When the request completes, the conversation would be returned to the pool. The second conversation is used and returned when the directory is created. The third conversation is used when file 1 is opened and returned when it is closed. The fourth conversation is used when file 2 is opened and returned when it is closed.

Now suppose that the application is changed in the following manner:

1. Open file 2.
2. Read from file 2.
3. Find the available space on a volume.
4. Create a directory
5. Open file 1.
6. Write to file 1.
7. Close file 2.
8. Close file 1.

These changes have an impact on the use of conversations.

A conversation is allocated when file 2 is opened, and continues to be used as long as there is at least one open file. This application uses only one conversation while the other requires four.

A disadvantage to this approach is that the conversation is not being shared. Assume additional processing is taking place between each of the listed steps. There would be periods of time where the conversation could be returned to the pool but is not because a file is open. If you had a large number of applications using a small pool of conversations, you may quickly run out of conversations. Jobs needing a conversation would go into a wait state until conversations become available.

Minimizing Volume Mounts for Optical LAN Server

Other factors that can greatly affect performance for both optical LAN servers and directly-attached libraries are volume mounting and drive contention. In order to remove and mount a new volume, the following sequence of events takes between 8 to 15 seconds.

- Remove a volume from the optical drive.
- Store the volume in an open library slot.
- Select the new volume to mount from a library slot.
- Move the volume into the optical drive.

Applications should be designed to minimize volume mounting. The optical LAN server also has its own set of performance parameters that can be changed to optimize optical drive usage. Refer to the *IBM 3995 LAN Optical Library Dataserver Reference* book for more information about the performance parameters available for LAN-attached optical media libraries.

Chapter 10. Optical Media Formats

OS/400 provides support for several optical media types and media formats. The **media format** is the file system architecture that exists on the media to manage file, directory, and volume information. The OS/400 supports the following optical media formats.

- ISO 9660

This industry-standard media format specifies the volume and file structures of compact read only optical disks (CD-ROM).

- HPOFS (High Performance Optical File System)

This IBM developed architecture is the media format that is used for WORM and erasable media that are supported in 3995 optical media library devices. The OS/400 implementation of HPOFS is the WORM (write-once-read-many) based version of this architecture. Everywhere in this chapter where the term **HPOFS** is used, the WORM version of HPOFS is implied.

- UDF (Universal Disk Format)

This industry-standard subset of ISO 13346 is the media format that is created on DVD-RAM media through OS/400. It is also available to use on erasable media in 3995 optical media library devices.

Writable optical media (WORM, erasable, DVD-RAM) is initialized on OS/400 using the Initialize Optical (INZOPT) CL command. WORM media must use media format HPOFS. DVD-RAM media must use media format UDF. Erasable media may use either HPOFS or UDF depending on the requirements of the user. You can specify the format by using the MEDFMT keyword on the INZOPT command.

This chapter provides information about the different media formats. This chapter also provides a comparison of the media formats so users of erasable media can select the media format that best meets their requirements.

ISO 9660

Overview

This industry standard media format was originally designed to specify the volume and file structures of compact read only optical disks (CD-ROM), and is a read-only media format. It is used currently on CD and DVD read only media. OS/400 supports ISO9660 media that is created using the primary volume descriptor (PVD). OS/400 does not support ISO9660 extensions that use the supplementary volume descriptor (SVD).

In addition, there is some support for ISO9660 extensions defined as the IEEE P1281 (System Use Sharing Protocol) and P1282 (Rock Ridge Interchange Protocol) specifications. These extensions are also known as Rock Ridge. OS/400 supports the Rock Ridge alternate name structures through the PVD. This support provides recognition and handling of mixed case and long file names, similar to UNIX. Other structures defined by Rock Ridge are not supported by OS/400. To enable support to read Rock Ridge alternate name structure (if it exists), enter CHGOPTA EXTMEDFMT(*YES) prior to importing the media. If the Rock Ridge alternate name structures do not exist, entering this command will have no effect.

Volume, Directory, and File Names

The volume identifier for the primary volume descriptor can be a maximum of 32 characters. The volume identifier must contain only alphabetic characters (A through Z), numeric characters (0 through 9), or the underscore (_).

Although not required, you can include one or more directories in the path name. Each element of the path can be a maximum of 32 characters with the total maximum path length of 256 characters. A path name can consist of any alphabetic characters (A through Z), numeric characters (0 through 9), or the underscore (_).

For ISO9660 media containing Rock Ridge extensions, the length of each element name is not restricted, but the total maximum path length is still 256 characters. The pathname characters are not restricted, but are recommended to consist of the POSIX portable filename character set ("A" through "Z", "a" through "z", "0" through "9", period (.), underscore (_), or hyphen (-)).

File name searches are case-insensitive meaning that you can use either uppercase or lowercase characters to access existing files.

For ISO9660 media containing Rock Ridge extensions, file searches are case-sensitive. If no case-sensitive match is found, a case-insensitive match is returned if it exists. If multiple case-insensitive matches exist on the volume, an error will be returned indicating that ambiguous names exist on the media. Some optical commands, such as Copy Optical (CPYOPT), are not supported when duplicate, ambiguous file names exist. For example Rock Ridge allows files "ABC.ext" and "abc.EXT" to exist in the same directory. This is not supported by CPYOPT and may produce unpredictable results.

Programming Interfaces

The system can read files on ISO 9660 media by using either the HFS (Hierarchical File System) programming interface or the integrated file system programming interface.

Command Interfaces

You can use OS/400 restore commands to restore data from ISO 9660 media. This is true only when the media was correctly mastered from a save image on tape.

There are some restrictions on which optical commands are supported to ISO 9660 media. Refer to Table 10 on page 120 later in this chapter for more information about which commands the system supports for each media format.

Directory and File Security

There is no directory and file level security for ISO 9660 media. Volume level security is available through authorization lists. See Chapter 6, "Optical Security and Auditing," on page 53 for more information.

High Performance Optical File System

Overview

HPOFS (High Performance Optical File System) is an IBM developed media format architecture available to use when initializing optical media on the OS/400. OS/400 uses the WORM-based version of HPOFS. This media format is designed (and required) for WORM media, but you can use it (and it is the default) when

initializing erasable optical media. HPOFS is a WORM media format. You can write each sector of the media only once when creating and updating files and directories. This unique characteristic of never rewriting the same sector allows all previous versions of every file to remain on the media. One drawback of this is that media consumption continues to grow as you update and even delete files.

This section contains detailed information on the OS/400 implementation of HPOFS for direct-attached optical media libraries. This section does not address HPOFS characteristics for LAN-attached optical media libraries.

Volume, Directory, and File Names

Volume identifiers can be a maximum of 32 characters and must contain only alphabetic characters (A through Z), numeric characters (0 through 9), a hyphen (-), or a period (.). The first character must be alphabetic or numeric, and the identifier cannot contain imbedded blanks.

Although not required, you can include one or more directories in the path name. Each element of the path can be a maximum of 255 characters with the total maximum path length of 256 characters. A path name can consist of any of the EBCDIC characters except x00-x3F, xFF, ", *, <, >, ?, and \.

The system stores all alphabetic characters for directory and file names to the media in uppercase. File name searches are case-insensitive meaning that you can use either uppercase or lowercase characters to access existing files.

Space Reclaim

You can update or delete files even though HPOFS is a WORM media format. When a file gets changed or deleted, a new version of the file gets written, and the old version still exists on the media. This is true for both WORM and erasable media. The old file versions will always exist on WORM and will exist on erasable until the entire volume is reinitialized. When you change or delete a file, the system does not reclaim the space that was used by the old file. Media consumption continues to increase on HPOFS media until you reinitialize the volume (for erasable media). You can never reclaim deleted space for WORM media.

Programming Interfaces

You can create or read files on HPOFS media by using either the HFS (Hierarchical File System) programming interface or the integrated file system programming interface. There are things specific to the OS/400 implementation of the HPOFS media format that application writers need to know.

- National Language Support

The integrated file system interfaces assume that the CCSID (coded character set identifier) represents the path in effect for the job. The system then converts the path from the job CCSID to an internally used CCSID. The HFS interface makes no assumption about the CCSID of the path; therefore the system performs no character set conversion on the path. This could produce undesirable side effects if an application used the two API sets interchangeably or if the application changed from using one API set to the other.

A program should not create files through HFS and then try to read them by using the integrated file system APIs. Depending on the characters that are used in the path name, this could result in a "File not found" error. This type of problem can be avoided if the application uses only invariant characters (for

example A-Z 0-9 + = % & () , _ . : ;) for path names. Invariant characters are graphic characters that map to the same code point in all character sets.

- Held optical files

When the system writes an optical file that it cannot close normally, the system may create a held optical file. The held file exists on OS/400 internal disk storage and contains the data written to the file. You can then save or release the held file through an API or command interface. The system creates held files only when files fail to archive on HPOFS media.

- Synchronous writes

You can open files on HPOFS through HFS by specifying all writes to be synchronous. When specified, writes will be synchronous to OS/400 internal disk storage, not to the optical media. In the event of a power failure the data is recoverable from a held optical file.

Similarly for the HFS Force Buffered Data API and the integrated file system fsync() API, data will be forced to OS/400 internal disk storage, not to optical. Again, in the event of a power failure, the data is recoverable from a held optical file.

- File sharing

Multiple jobs or threads can share files. The system fully honors file that share modes as specified on the open request. For example, assume that a job opens a file that specifies it to share only with readers. This means that you can perform other opens only as long as the access requested remains Read Only.

- Extended file attributes

The system supports extended file attributes for files on HPOFS media. Extended attributes can be written to files using the HFS Change Directory Entry Attributes API as well as through some integrated file system interfaces.

Command Interfaces

OS/400 save and restore commands can be used to save and restore data on HPOFS optical media. Refer to Appendix B, "Optical Save and Restore," on page 129 for more information about save and restore to HPOFS volumes.

There are no restrictions on which optical commands the system supports for the HPOFS media. Refer to Table 10 on page 120 later in this chapter for more information about which commands the system supports for each media format.

Directory and File Security

There is no directory and file level security for HPOFS media. Volume level security is available through authorization lists. See Chapter 6, "Optical Security and Auditing," on page 53 for more information.

Media Interchange

HPOFS formatted optical media created in a direct-attached optical library can be accessed in a LAN-attached optical library assuming that the media type is supported. Conversely, optical media created in a LAN-attached optical library can be accessed in a direct-attached if the media type is WORM or is Erasable and formatted using "WORM Format" that does not support the mixed uppercase and lowercase structure.

UDF formatted optical media created in a direct-attached optical library cannot be accessed in a LAN-attached optical library device.

Directory Structure and Performance

HPOFS volumes have a dual directory structure to access files. Both a hash and hierarchical structure exist to provide a primary and secondary path to the file data. If the primary directory structure becomes damaged, the secondary is used.

The hash directory structure is designed to reduce the amount of media I/O required which improves performance for file access. Because of this hash directory structure, directory depth has less effect on performance than if the directory were searched hierarchically. For example, if /DIRECTORY1 contains 1000 files and /DIRECTORY2 contains 100 files; file search times for files in /DIRECTORY1 will generally take no longer than file searches in /DIRECTORY2. This is because the system performs the searches by using the hash structure, not the hierarchical structure.

Directory depth has less effect on performance for a hash search than for a hierarchical search. However, the overall directory depths and total number of files on a volume will effect performance. In general, a volume with fewer files on it will result in better file performance than a volume with more files.

Universal Disk Format

Overview

UDF (Universal Disk Format) is the OSTA (Optical Storage Technology Association) supported subset of ISO/IEC 13346. It also addresses ECMA-167 which is equivalent to ISO 13346. UDF is a writable file format that provides true space reclaim capabilities as well as file and directory level security. This section contains detailed information on the OS/400 implementation of UDF for direct-attached (C4x) 3995 optical media libraries as well as a DVD-RAM device.

Volume, Directory, and File Names

Volume identifiers can be a maximum of 30 characters and must contain only alphabetic characters (A through Z), numeric characters (0 through 9), a hyphen (-), or a period (.). The first character must be alphabetic or numeric, and the identifier cannot contain imbedded blanks.

Although not required, you can include one or more directories in the path name. Each element of the path can be a maximum of 254 characters with the total maximum path length of 256 characters. A path name can consist of any of the EBCDIC characters except x00-x3F, xFF, ", *, <, >, ?, and \.

The system stores all alphabetic characters for directory and file names to the media in uppercase when created through HFS or the OS/400 save interfaces. The system stores all alphabetic characters for directory and file names to the media in mixed case when created through the integrated file system interfaces. File name searches are case-insensitive meaning that you can use either uppercase or lowercase characters to access existing files.

File searches on UDF volumes created by OS/400 are case-insensitive. For UDF media created or updated by another operating system platform, a case sensitive search is performed. If no case-sensitive match is found, a case-insensitive match is returned if it exists. If multiple case-insensitive matches exist on the UDF volume, an error will be returned indicating that ambiguous names exist on the media. Some optical commands, such as Copy Optical (CPYOPT), are not supported when duplicate, ambiguous file names exist. For example UDF created on another

| operating system may allow files "ABC.ext" and "abc.EXT" to exist in the same
| directory. This is not supported by CPYOPT and may produce unpredictable
| results.

Programming Interfaces

You can create files or read files on UDF media by using either the HFS (Hierarchical File System) programming interface or the integrated file system programming interface. There are things specific to the OS/400 implementation of the UDF media format that application writers need to know.

- National Language Support

The integrated file system interfaces assume that the CCSID (coded character set identifier) represents the path in effect for the job. The system then converts the path from the job CCSID to an internally used CCSID. The HFS interface makes no assumption about the CCSID of the path; therefore the system performs no character set conversion on the path. This could produce undesirable side effects if an application used the two API sets interchangeably or if the application changed from using one API set to the other.

You should not create files through HFS and then try to read them by using the integrated file system APIs. Depending on the characters that are used in the path name, a "File not found" error could result. This type of problem can be avoided if the application uses only invariant characters (for example A-Z 0-9 + = % & () , _ . : ;) for path names. Invariant characters are graphic characters that map to the same code point in all character sets.

Since UDF is an industry-standard media format, NLS compliance could be important because of the increased opportunity for media interchange across different operating system platforms. This causes the system to limit HFS interfaces to UDF media. Additionally, the system uses invariant characters to reduce the chance of media interchange problems that are related to file names. Assume that an HFS-based application absolutely requires the use of variant characters. You can use the Change Optical Attributes (CHGOPTA) CL command to allow variant characters through the HFS interface by specifying CHGOPTA ALWVRNT(*YES). Once the system allows variant characters through HFS, there is no guarantee that path names will interchange correctly if accessed from another operating system. There is also no guarantee that path names will be consistent between the HFS and integrated file system interfaces.

- Held optical files

The system does not create held files for UDF media. When a file fails to close on UDF, the system signals an error to the application. This error then closes the file without writing the data to optical disk. The application must rewrite the file (open, write, close) to ensure that the data is on optical disk. The exception to this is if the application did a Force Buffered Data or fsync() prior to the close. These APIs will force the writing of the data to the optical disk.

- File sharing

Multiple jobs or thread can share files for read, but writers are always exclusive. Remember: If one job or thread is writing to a file on UDF, you cannot use any other jobs or threads to open that file.

Therefore, when using integrated file system open() or open64() APIs, the sharing modes O_SHARE_RDONLY, O_SHARE_WRONLY, and O_SHARE_RDWR do not provide the requested level of sharing when the access mode is O_RDWR or O_WRONLY. When the access method is O_RDWR or O_WRONLY, the resulting sharing mode will be equivalent to O_SHARE_NONE.

When using the HFS Open Stream File API, the lock modes Deny None, Deny Write, and Deny Read do not provide the requested level of sharing when the access mode is Write Only or Read/Write. When the access method is Write Only or Read/Write, the resulting lock mode will be Deny Read/Write.

- Mixed-case file name

When created through the integrated file system interfaces, files and directories created on UDF volumes will preserve the case specified on the create. For example, if file "Abc" is specified on the open() API, "Abc" will be created on the media in the mixed-case form. Even though the system preserves file case, file searches are case-insensitive meaning that the system can read the file that uses any case such as "ABC" or "abc".

When created through the HFS or save and restore interfaces, the system stores files and directories that are created on UDF volumes in uppercase. For example, if you specify file "Abc" on the Open Stream File API, the system creates "ABC" on the media. Again, file searches are case-insensitive so you can specify any case to read the file.

File searches on UDF volumes created by OS/400 are case-insensitive. For UDF media created or updated by another operating system platform, a case sensitive search is performed. If no case-sensitive match is found, a case-insensitive match is returned if it exists. If multiple case-insensitive matches exist on the UDF volume, an error will be returned indicating that ambiguous names exist on the media.

Command Interfaces

OS/400 save and restore commands can be used to save and restore data on UDF optical media. Refer to Appendix B, "Optical Save and Restore," on page 129 for more information about save and restore to UDF volumes.

There are some restrictions on which optical commands are supported to UDF volumes. For example, the system does not support Work with Optical Directories (WRKOPTDIR) and Work with Optical Files (WRKOPTF). Display Optical (DSPOPT) has some restrictions when used for UDF volumes. You should use the integrated file system commands Work with Object Links (WRKLNK) and Display Object Links (DSPLNK) instead of the optical commands. Refer to Table 10 on page 120 later in this chapter for more information about which commands the system supports for each media format.

Directory and File Security

Directory and file level security is available for UDF volumes. The system maintains the "data authorities" of optical directories and files for three groups of users; they are: owner, group, and public. Volume level security is also available through authorization lists. See Chapter 6, "Optical Security and Auditing," on page 53 for more information.

Directory and file level security is not guaranteed when volumes are removed and transported to and from other systems. Security information recorded in the UDF structures on the media may not have the same meaning on another system as it does on the system where it was written.

Media Interchange

UDF media created on OS/400 is UDF Version 2.01. This media will interchange to other operating system platforms that support this version of UDF.

UDF compliant media that is created with UDF Version 1.5 or less will be accessible by OS/400 as "read only". Media created with UDF Version 2.0 and UDF 2.01 will be accessible by OS/400 for read and write.

Directory Structure and Performance

UDF volumes have a single (hierarchical) directory structure to access files. Because of this hierarchical directory structure, the depth of a directory tree has a direct impact on file performance. For example if /DIRECTORY1 contains 1000 files and /DIRECTORY2 contains 100 files, file search times for files in /DIRECTORY1 will, in general, take longer than file searches in /DIRECTORY2. This is because the system performs file searches hierarchically, which may require looking at every entry in the directory.

In general, file performance will be better for UDF if you evenly distribute files across several directories and subdirectories.

Command and Media Format Dependencies

Some optical commands have no meaning when used with certain optical media formats. No support exists for other commands with certain optical media formats. Table 10 lists all of the volume related optical commands and the media formats to which they apply.

Table 10. Optical Commands and Media Format Dependencies

Command	ISO 9660	Universal Disk Format	HPOFS in Directly-Attached Device	HPOFS in LAN-Attached Device
CHGOPTVOL	Partially supported ¹	Partially supported	Supported	Partially supported
CHKOPTVOL	Supported	Supported	Supported	Not supported
CPYOPT	Supported	Supported	Supported	Not supported
CVTOPTBKU	Not applicable	Not applicable	Supported	Not applicable
DSPOPT	Supported	Partially Supported	Partially supported	Partially Supported
DSPOPTLCK	Supported	Supported	Supported	Partially supported
DUPOPT	Not supported	Supported	Supported	Not supported
INZOPT	Not applicable	Supported	Supported	Not supported
WRKHLDOPTF	Not applicable	Not applicable	Supported	Not applicable
WRKOPTDIR	Supported	Not supported	Supported	Partially supported
WRKOPTF	Supported	Not supported	Supported	Partially supported
WRKOPTVOL	Supported	Supported	Supported	Partially supported
Notes:				
1. <i>Partially supported</i> indicates that not all command parameters may apply when used with the indicated device.				

Differences Between HPOFS and UDF

There are some significant differences between the OS/400 implementation for HPOFS and UDF. It is important to understand these differences when selecting a media format when initializing your erasable optical media. It is also important if you are an application writer using either HFS or the integrated file system to store and retrieve files from optical.

Here are the major differences between the two media formats:

- Media format and media type
 - HPOFS is available on *WORM or *ERASE media types.
 - UDF is available on *DVD-RAM and *ERASE media types.
 - HPOFS is the default media format for *ERASE media type.
- Space Reclaim

The system reuses or reclaims media space on UDF media when it deletes or updates files. The system does not reuse media space on HPOFS when it deletes or updates files. For HPOFS on *ERASE media type, space can be reclaimed through "bulk erase" by reinitializing the volume.
- File and directory authority

File and directory-level authority (permissions) is available on UDF volumes. This level of authority is not available for HPOFS files.
- Command support

The system does not support several optical commands such as WRKOPTDIR and WRKOPTF when using UDF volumes.
- Volume id and path names
 - HPOFS volume ids can be 32 characters long. UDF volume ids can be 30 characters long.
 - A single path element can be a maximum of 255 characters for HPOFS and 254 characters for UDF.
 - In order to use variant characters to UDF through HFS, you must use the Change Optical Attributes (CHGOPTA) CL command as follows:
`CHGOPTA ALWVRNT(*YES)`
 - The system preserves file case when creating files on UDF through the integrated file system interfaces. The system always stores file names in uppercase on HPOFS and UDF volumes when creating them through HFS and Save and Restore.
- Writers to UDF files are exclusive - no sharing with writers.
- The system does not create held files for UDF files.
- Synchronous writes

Force Buffered Data and fsync() to HPOFS will force the data to OS/400 internal disk storage. Recovery is through a held optical file. When you issue these APIs to files on UDF volumes, the system writes the data to optical disk.
- Save and Restore

Save and Restore multi-volume processing works differently for UDF than it does for HPOFS. See Appendix B, "Optical Save and Restore," on page 129 for more information.
- File performance

Since UDF is a hierarchical file system, performance is directly affected by the depth of the directory tree on the volume. For best results, you should evenly

distribute files across directories on the volume. With the hash directory structure on HPOFS, the depth of the directory tree has less impact on performance than it does for UDF.

Appendix A. Reclaiming the Optical Index Database

A system level index, called the **optical index database**, keeps track of all optical volumes and directories known to the system. This database includes the optical volume index (QAMOVAR) and the optical directory index (QAMOPVR) physical files. You can use the Reclaim Optical (RCLOPT) command to re-create the optical index database if it is ever damaged or destroyed or whenever volumes that you know are in an optical media library, CD-ROM, or DVD device are reported as not found. To run the RCLOPT command, either select option 2 (Reclaim optical index) on the Optical Backup/Recovery display or enter the RCLOPT command. Doing either causes the Reclaim Optical (RCLOPT) display to appear.

Note: The RCLOPT command (shipped with a public authority of *EXCLUDE) applies to directly-attached optical media libraries, CD-ROM, and DVD optical devices. You cannot issue the RCLOPT command to LAN-attached optical media libraries. To re-create the Optical Index Database for LAN-attached optical media libraries, use the Add Optical Server (ADDOPTSVR) command.

Optical Index Information

Optical index information regarding which volumes are in a particular optical media library and which directories are on each volume is kept in optical index files at different levels within the system. These files are used to enhance performance by eliminating the need to access the optical media library or physical media each time the location of a volume or directory is needed.

Failures, system upgrades, and physically moving optical library devices from one system to another can cause these index files to become out of synchronization with the actual contents of a particular optical media library or volume. When this happens, messages are sent indicating that the optical index needs to be reclaimed, such as OPT1245, OPT1825, or OPT1330. These messages direct you to run the RCLOPT command. The following topics describe the optical index files that are kept at the different levels of the system. An understanding of the different optical indexes is helpful when deciding which type of reclaim optical index to run.

Optical Index Database Files

The optical indexes reside in the physical files QAMOVAR and QAMOPVR. The QAMOVAR file is the optical volume index. It contains information about all optical volumes known to the system. This includes volumes that were previously removed from the optical media library with the volume description option of *KEEP. The QAMOPVR file is the optical directory index. It contains information about the directories on the volumes in directly-attached optical media libraries or CD-ROM devices. This includes those volumes that were previously removed from directly-attached libraries with the volume description option of *KEEP.

Information for volumes that are *OFFLINE or *REMOVED is retained by reclaim optical processing, but it cannot be rebuilt or verified because the physical volumes are no longer accessible. If the optical index database is ever destroyed, information about *REMOVED volumes can be recovered by adding the cartridge that contains the volumes to an optical media library using the Add Optical Cartridge (ADDOPTCTG) command.

Internal Library Index

Each optical media library keeps an internal library index of each volume that it contains. The internal library index for each optical media library is controlled by the Licensed Internal Code. The information in this index is generally not accessible to users or application programs. However, this index must be kept synchronized with the optical index database. This index is re-created when the *RESET rebuild type is specified.

To select the optical media library or libraries that require rebuilding, enter the optical media library name in the *Optical media library* field on the Reclaim Optical (RCLOPT) display. The name that you enter must correspond to an optical media library that is currently configured on the system.

To reclaim more than one optical media library, issue separate RCLOPT commands for each device rather than using MLB(*ALL). Sequentially using MLB(*ALL) reclaims all of the optical media libraries one at a time. Issuing separate RCLOPT commands will allow reclaims to run in parallel which will run faster.

Reclaiming the Optical Index for a Stand-Alone Optical Device

The optical index database re-creates entries for CD-ROM and DVD devices each time the device is varied on with media in the device. The easiest way to reclaim the optical index for a stand-alone optical device is to vary the device description off and on again using the Vary Configuration (VRYCFG) command. Ejecting and re-inserting the media has the same effect. You can issue the RCLOPT command for stand-alone optical devices if you choose to do so.

Reclaim Types

There are three possible types to select: *SYNC, *UPDATE, and *RESET. Each successive reclaim type described in the following topics is more extensive and takes longer to run. The *UPDATE and *RESET reclaim types allow you to optionally reclaim the optical directory index. The type of index problem that you are experiencing determines which reclaim option should be run. See “Choosing the Reclaim Type to Use” on page 126 for more information about determining which option to use. The following topics describe the processing during each reclaim option.

***SYNC—Synchronize Volume Index with Internal Library Index**

The synchronize option verifies that the entries in the optical index database are also in the internal library index. Entries that are in both indexes are left unchanged. Only those optical volumes that are in the internal library index but not in the optical index database are mounted in an optical drive.

If an entry is in the internal library index but not in the optical volume index, an entry is created for the volume in the optical volume index. Message OPT2105 (Optical index entries created for volume &2;) is issued, indicating that an optical volume index entry was created for the volume. If the volume is initialized, optical directory index entries are also created for each directory on the volume.

If an entry is in the optical volume index but not in the internal library index, message OPT2115 (Optical volume &1; is marked removed) is issued. This indicates that the volume status for that volume has been changed to *REMOVED.

The *SYNC reclaim type requires exclusive use of all libraries whose indexes are being reclaimed. Also, when using the *SYNC reclaim type, the Volume identifier and Reclaim directory index fields are not used.

***UPDATE—Update Volume Index from Optical Volumes**

The update option re-creates the optical volume index entries for all volumes in a media library or a specific volume by reading the volume data from the media. In addition, you can optionally rebuild the optical directory index by using the DIR parameter.

If *ALL is specified in the Volume identifier field, the optical volume index is reclaimed for all volumes in the optical media library. If a specific volume name is entered in the Volume identifier field, the optical volume index is reclaimed for that volume only. The optical volume index is updated only for those volumes and libraries that are selected. Index information for other volumes and libraries remains unchanged. Each optical volume whose index is reclaimed will be mounted in an optical drive.

If all of the volumes in an optical media library are specified and an entry is in the internal library index but not in the optical volume index, an entry is created for the volume in the optical volume index and message OPT2105 is issued. If the volume is initialized, optical directory index entries are also created for each directory on the volume.

If an entry is in the optical volume index but not in the internal library index, message OPT2115 is issued. This message indicates that the volume status for that volume has been changed to *REMOVED.

Reclaim Optical Directory Index Option

The update option lets you reclaim the optical directory index (QAMOPVR) file. The following values are available:

- *YES indicates that the optical directory index is reclaimed for each volume whose optical volume index is reclaimed.
- *NO indicates that the optical directory index is not reclaimed for the volume.

The *UPDATE type requires exclusive use of all volumes that are being updated. If an optical media library of *ALL is specified, this reclaim type requires exclusive use of all libraries whose indexes are being reclaimed.

***RESET—Reset Internal Library Index and Reclaim Volume Index**

The reset option performs basically the same processing as the update option, except that the internal library index is reclaimed before the optical volume index is reclaimed. You can request that the internal library index and optical index database be re-created or updated either for a specific optical media library or for all optical media libraries. The optical volume index is updated only for those libraries that are selected. Index information for other libraries remains unchanged.

Specifying the *RESET option will always reclaim the optical directory index. The *RESET option requires the mounting of each cartridge in the optical media library at least once. The system does this to verify that the internal library index is correct.

If an entry is in the internal library index but not in the optical volume index, the system mounts and reads the volume again. The system creates an entry for the volume in the optical volume index and issues message OPT2105.

If an entry is in the optical volume index but not in the internal library index, the system issues message OPT2115. This message indicates that volume was not located after the internal library index rebuild, and that the volume status for that volume is changed to *REMOVED.

Reclaim Optical Directory Index Option

The reset option lets you reclaim the optical directory index (QAMOPVR) file. The following values are available for parameter DIR:

- *YES indicates that the optical directory index is reclaimed for each volume in the specified library.
- *NO indicates that the system does not reclaim the optical directory index for the volume.

The *RESET type requires exclusive use of all libraries that are being reclaimed. Also, when you use the *RESET type, the Volume identifier field is not used.

Time Required to Complete Reclaim Optical Index

When either *RESET and VOLUME(*ALL) or *UPDATE and VOLUME(*ALL) together are selected, it may require hours for the Reclaim Optical Index command to complete. This is because every volume in the optical media library that is specified must be mounted and then read. The requested databases are then updated before the next volume is mounted. The following factors affect how long it takes the RCLOPT command to complete:

- How many libraries are being reclaimed
- How many volumes are in each library
- What type of reclaim is requested
- How many directories are on each volume

Once a reclaim command has started, it should not be canceled before it has completed. If a reclaim command is canceled before it completes, it might be necessary to run the reclaim command again before the optical media library is in a usable state.

Choosing the Reclaim Type to Use

Most optical support error messages that direct you to run reclaim optical index specify the rebuild type you should use to recover from the error. However, occasionally you might suspect that the optical index needs to be reclaimed even though no error message has been issued. In this case, you need to determine which reclaim type should be run.

If you are unsure which reclaim type you should use, run the RCLOPT command with the *SYNC option and then try the failing request again. If the request still fails, run the RCLOPT command with the *RESET option.

Following are hints on when to use the different reclaim options:

***SYNC**

Use this option when you are getting messages indicating that a volume is not found (OPT1331, OPT1330 - reason code 2) or that a volume is removed (OPT1460) when you feel that the volume is indeed in the optical

media library. Use this option after you upgrade to a new release of OS/400 or when you move a directly-attached optical library device from one server to another.

***UPDATE**

Use this option first if you see a message indicating that the optical tables are incorrect (OPT1825). You can also use this option if you are having problems with a particular volume not displaying all the directories when you use the Work with Optical Directories (WRKOPTDIR) command.

***RESET**

Use this option when you get a message OPT1330 with reason code 01. Unless otherwise instructed through an optical message, use this option as a last resort. It will generally take much longer to complete than either of the two previous options, but it will ensure that both the optical index database and the internal library index are correct. Specify DIR(*NO) unless you have a specific need to create the optical directory index. The only operations that require the directory index are Work with Optical Directories (WRKOPTDIR) and Display Optical (DSPOPT) when DATA(*DIRATR) is used. If you specify DIR(*NO), the directory index will be built on demand when one of these functions is issued.

Appendix B. Optical Save and Restore

Introduction

The OS/400 Save and Restore CL commands support directly-attached optical library devices, CD-ROM, DVD-ROM, and DVD-RAM standalone devices.

The best use of optical storage devices is as a part of disaster recovery protection. The extraordinary long shelf life of optical media is well suited for the long term storage of critical data. You can provide extra protection by using permanent WORM media, because you cannot alter data on the media.

Tape devices may provide the optimal day to day backup mechanism. This depends on the amount of data you want backed up, and the amount of system time available for backup.

CD-ROM and DVD-RAM media are also very well suited for software distribution. The save/restore command interface can be used as a part of installation procedures for programs, data, and program fixes. CD-ROM and DVD-RAM standalone drive optical devices also support the LODRUN CL command.

Save and Restore Command Summary

Optical devices support many of the most widely used OS/400 Save and Restore CL commands.

The DVD-RAM standalone optical drive device is an economical alternative to magnetic tape for save/restore operations on entry level iSeries servers. **The DVD-RAM standalone drive device supports all major save/restore functions.**

Refer to “Command Restrictions” on page 134 for a list of the save and restore commands restrictions. The automated library device enhances ease of use of save and restore operations that require volume lists.

See the “Command Restrictions” on page 134 section for more information on command operation.

Refer to the CL topic in the Programming category of the Information Center for detailed information on Save and Restore command syntax and functions.

Operational Hints

Optical Media Formats

| Save and Restore volume list processing differs for each optical media format. A
| **volume list** is used on a save or restore request when multiple optical volumes are
| required for the operation, thus creating a **volume set**. All volumes in a volume set
| must have the same optical media format. Volume sets are not supported for CD
| media formatted with ISO 9660.

| You must initialize media of type Permanent WORM and CCW WORM with a
| media format of HPOFS.

You can initialize media of type Rewritable with a media format of HPOFS or UDF (Universal Disk Format).

You must initialize the DVD type of media that is used by DVD-RAM standalone drive devices with a media format of UDF (Universal Disk Format).

Organizing Save and Restore Data on Optical Devices

The save data on optical media is uniquely identified by a pathname. This pathname has the form: `/directoryname/subdirectoryname/./filename`. You may create and specify as many directory levels as is necessary to organize your save data to suit your needs. If no directory levels are specified, the save data file is placed in the root directory of the specified optical volume.

The optical pathname may be up to two hundred and fifty-six alphanumeric characters in length. Optical volume names may be up to thirty-two alphanumeric characters in length. Some caution is necessary when employing long names. Many OS/400 save and restore displays, messages, reports, output files, and object descriptions, support a maximum of six characters for volume names, and seventeen characters for pathnames. Longer names will appear truncated in these instances. Additionally, some automated data management software may not properly handle long volume names and long pathnames.

Performing a Save Operation to Optical Storage

For example, you can save the OS/400 library "DEVLIB01" to the optical volume "SRVOL1" that is contained in library device "OPTMLB02". You do this by using the following CL command:

```
SAVLIB LIB(DEVLIB01) DEV(OPTMLB02) VOL(SRVOL1) ('/DEVLIB01')
```

An optical file containing the save data, with the name DEVLIB01, will be created in the root directory of volume SRVOL1.

Displaying Save and Restore File Information on Optical Storage

As an example, information concerning the save and restore files that are contained on a given optical volume can be displayed using the **DSPOPT CL** command. The following CL command displays the information for all save and restore files that are found in the root directory of the optical volume SRVOL1:

```
DSPOPT VOL(SRVOL1) DATA(*SAVRST) PATH(/)
```

OPTFILE Parameter

The OPTFILE parameter is used in Save and Restore commands to designate the optical file pathname to be used to contain the save data.

The system dynamically creates any specified directory names that do not exist.

The OPTFILE parameter has a default value of `""`. Using the default parameter value places the file in the root directory of the optical volume that is specified by the VOLUME parameter. Additionally, in commands other than **SAV**, the file name is the name of the OS/400 library that contains the objects that are saved.

For the **SAV** command, **OPTFILE(*)** generates a filename of the form: `SAVyyyymmddhhmmssmmm`, where `yyyymmddhhmmssmmm` is the current date and time.

Media Eject Option

For standalone drive devices that are attached by *PowerPC* I/O adapters, you can automatically open the media tray at the conclusion of a Save and Restore operation. You can do this by specifying the **ENDOPT(*UNLOAD)** parameter. The system ignores this parameter for optical library devices. The **ENDOPT(*LEAVE)** or **ENDOPT(*REWIND)** parameters have no effect for optical standalone drive devices or optical library devices.

Volume List Processing

Volume lists allow a single Save and Restore operation to use many pieces of optical media to complete the requested operation. Information relating to optical volumes that are a part of a save and restore volume list may be display using the **DSPOPT** command.

Volume list information fields:

- **Continued from Previous Volume** Flag
 - This field applies only for UDF media
 - Indicates the save and restore file is continued from the previous volume in the volume list
 - Only one file on a given volume can have this flag set on.
 - Use the **DSPOPT DATA(*FILATR)** command to display this flag.
- **Continued on Next Volume** Flag
 - This field applies only for UDF media.
 - Indicates the save and restore file is continued to the next volume in the volume list.
 - Only one file on a given volume can have this flag set on.
 - Use the **DSPOPT DATA(*FILATR)** command to display this flag.
- **IPL Capable** Flag
 - This flag indicates that the SAVSYS command created this volume and that you can use it for D-Mode IPL.
 - Use the **DSPOPT DATA(*VOLATR)** command to display this flag.
- **Last Volume in Volume List** Flag
 - This flag indicates that the volume is the final volume in a volume list.
 - For HPOFS format volumes, the system does not allow save files unrelated to the volume list on the final volume. The remaining capacity does not affect this situation. UDF format volumes will allow unrelated save files on the volume if sufficient free space exists.
 - Use the **DSPOPT DATA(*VOLATR)** command to display this flag.
- **Starting Volume ID** Field
 - Volume ID of the first volume in a multiple volume set (a volume list). For UDF format volumes, the volume list may contain several different save files. Consequently, this field does not specify the starting volume of any given file that is contained in the set.
 - Use the **DSPOPT DATA(*VOLATR)** or **DATA(*FILATR)** command to display this flag.

Volume Lists with HPOFS Format Media

Only one file in a multi-volume set logically "spans" volumes. The last volume in the set does not accept additional save requests. The system does not maintain "continued" flags.

- Volume1 (Sequence#=1, Starting valid=Volume1, Last volume in set=No)
 - File1 (Continued from previous volume=NO, Continued on next volume=NO)
 - File2 (Continued from previous volume=NO, Continued on next volume=NO)
 - File3 (Continued from previous volume=NO, Continued on next volume=NO)
- Volume2 (Sequence#=2, Starting valid=Volume1, Last volume in set=No)
 - File3 (Continued from previous volume=NO, Continued on next volume=NO)
- Volume3 (Sequence#=3, Starting valid=Volume1, Last volume in set=Yes)
 - File3 (Continued from previous volume=NO, Continued on next volume=NO)

Note: No more saves are allowed to Volume1, Volume2, or Volume3. The system will not allow additional save files to Volume3, regardless of sufficient free space.

Note: Access to any previous save data is lost from Volume1, Volume2, and Volume3.

Note: All restores must begin on Volume1.

Volume Lists with UDF Format Media

More than one file in a multi-volume set can logically "span" volumes, but only one file per volume can span to the next volume in the volume list. The last volume in the set will accept additional save requests if space is available. The system maintains the "continued" flags for spanned files.

- Volume1 (Sequence#=1, Starting valid=Volume1, Last volume in set=No)
 - File1 (Continued from previous volume=NO, Continued on next volume=NO)
 - File2 (Continued from previous volume=NO, Continued on next volume=NO)
 - File3 (Continued from previous volume=NO, Continued on next volume=YES)
- Volume2 (Sequence#=2, Starting valid=Volume1, Last volume in set=No)
 - File3 (Continued from previous volume=YES, Continued on next volume=NO)
 - File4 (Continued from previous volume=NO, Continued on next volume=NO)
 - File5 (Continued from previous volume=NO, Continued on next volume=YES)
- Volume3 (Sequence#=3, Starting valid=Volume1, Last volume in set=No)
 - File5 (Continued from previous volume=YES, Continued on next volume=YES)
- Volume4 (Sequence#=4, Starting valid=Volume1, Last volume in set=Yes)
 - File5 (Continued from previous volume=YES, Continued on next volume=NO)
 - File6 (Continued from previous volume=NO, Continued on next volume=NO)

Note: Space permitting, additional saves are allowed to Volume4.

Note: A restore begins on the volume that contains the first occurrence of the specified file. For example, you can restore data from File4 on Volume2 without processing Volume1.

Note: The "volume set" for UDF formatted media described in this section is not related to the "volume set" or "multi-volume support" that is defined in format specifications for UDF (ECMA 167 3rd Edition or UDF Revision 2.01). Support for volume sets as defined by these specifications is not currently provided by OS/400.

Performing a Save Operation to DVD-RAM, UDF, and HPOFS Media

Specifying the path for your files

Optical operates in random mode and uses a hierarchical file structure when writing files to the media. Beginning with the root directory of the volume, you may specify a path name for the optical file that is used for the save operation. Specifying '*' causes the system to generate an optical file name in the root directory '/'. Specifying 'optical_directory_path_name/*', causes the system to generate an optical file name in the specified directory of the optical volume. Specifying 'optical_file_path_name' creates an optical file name. For example, specifying SAVLIB LIB(MYLIB) DEV(OPT01) OPTFILE('/mydir/*') creates an optical file name of mydir/MYLIB. If mydir directory does not exist, the system creates it.

When you use DVD-RAM media to save OS/400 information, the system checks for active files by using the CLEAR() parameter on the save commands. Specify CLEAR(*NONE) to have the system search the DVD-RAM volume for any active optical files that have the same name. If an optical file of the same name exists, the system displays an inquiry message. You may cancel the processing, write over the existing file on the volume or insert a new cartridge. If no active files of the specified optical file exist and there is available space on your DVD-RAM volume, the system writes the file to your DVD-RAM media. If the system cannot find any available space on the media, it prompts you to insert a DVD-RAM volume in the device.

Using the CLEAR parameter

Specifying CLEAR(*ALL) will clear all the files on the media.

Specifying CLEAR(*AFTER) automatically clears all media after the first volume. The system sends an inquiry message when it encounters the specified optical file on the first volume. This allows you to either end the save operation or replace the file.

Specifying CLEAR(*REPLACE) automatically replaces the active data of the specified optical file on the media.

The only option on the CLEAR parameter that will clear all the files is CLEAR(*ALL). Otherwise, the system sends an inquiry message for each specified optical file name it encounters. Specifying CLEAR(*NONE) will send an inquiry message for each specified optical file name it encounters, not just the first one.

To avoid receiving an inquiry message during the save operation, you can perform either of the two following functions:

- Initialize the optical volume (INZOPT) first.

- Specify an option by using the CLEAR parameter on the save command.

Note: Do not use the CLEAR(*NONE) parameter, this sends an inquiry message.

IBM provides online information on using the CLEAR parameter with the save commands. Refer to the APIs topic in the Programming category of the Information Center for more information.

Using multiple volumes

After the system writes to a second DVD-RAM volume, the system considers this a DVD set. A set includes two or more volumes. The system can only write information on the last volume in the set. In a three volume DVD-RAM set, the system cannot write information to the first or second volume. See “Volume Lists with UDF Format Media” on page 132 for more information.

Command Restrictions

Command	Supported by Optical Device
SAVSTG	None
SAVS36F	None
SAVS36LIBM	None
SAVUSFCNR	None
RSTS36F	None
RSTS36FLR	None
RSTS36LIBM	None
RSTUSFCNR	None
SAVLICPGM	DVD (No optical libraries)
SAVSYS	DVD (No optical libraries)
RSTLICPGM	DVD (No optical libraries)
SAVCHGOBJ of more than one library, including LIB(*ALLUSR)	DVD and optical library devices with UDF formatted media
SAVDLO of more than one ASP	DVD and optical library devices with UDF formatted media
SAVLIB of more than one library, including LIB(*ALLUSR), LIB(*IBM), and LIB(*NONSYS)	DVD and optical library devices with UDF formatted media
SAVCFG	All writable optical devices
SAVCHGOBJ of one library	All writable optical devices
SAVDLO of one ASP	All writable optical devices
SAVLIB of one library	All writable optical devices
SAVOBJ of one library	All writable optical devices
SAVSAVFDTA	All writable optical devices
SAVSECDTA	All writable optical devices
RSTCFG	All optical devices
RSTLIB	All optical devices
RSTOBJ	All optical devices
RSTUSRPRF	All optical devices

Command	Supported by Optical Device
SAVAPARDTA	Not applicable, command does not use a device
RSTAUT	Not applicable, command does not use a device

Operational Differences by Device Types

Optical Library Dataservers

- You cannot use the default value, *MOUNTED, for the volume identifier.
- Volumes that are provided in a volume list must all be in the same library device.
- A single save data file may span several volumes in a volume list.
- For HPOFS formatted media, any volume used in a volume list, becomes unusable by any save or restore operation other than the operation originally processing the volume list.

For example:

- Save command A writes save data fileA to volume volA.
- Save command B writes save data fileB to volume list: volC, volB, volA.
- Restore command A **will not** be able to restore from fileA on volume volA.
- Restore command B **will** be able to restore from fileB on volume list: volC, volB, volA.

CD-ROM and DVD-ROM Standalone Optical Drive Devices

- CD-ROM and DVD-ROM are **read-only** devices. The system does not support Save commands for these devices.
- Save files can not span multiple CD-ROM or DVD-ROM media which contain ISO9660 media format.
- You can specify the default value, *MOUNTED, for the volume identifier. It will process the optical volume currently in the specified stand-alone device.

DVD-RAM Standalone Optical Drive Devices

- DVD-RAM devices are read and write devices. Save and restore commands are supported for DVD-RAM devices.
- You can specify the default value, *MOUNTED, for the volume identifier. It will process the optical volume currently in the specified stand-alone device.
- Multiple save data files may span several volumes in a specified DVD_RAM volume list.

Note: Software compression and decompression may increase the save and restore times. It uses considerable processing resources which may affect the overall system performance.

Appendix C. Recovering Held Optical Files

A **held optical file** is an optical file that could not be closed normally. It contains buffered data that cannot be written to the optical disk. If the open file handle is still valid, the file is still open; otherwise, it is considered closed. This appendix provides information on recovering held optical files.

Held optical files are only created for media format *HPOFS when the Change Optical Attributes (CHGOPTA) "Held file attribute field" is *YES. Held optical files are not created for media format *UDF or when the "Held file attribute" field for CHGOPTA is *NO.

How Optical Files Are Manipulated

An application can manipulate optical file data by using UNIX-type APIs or the hierarchical file system (HFS) as described in the Information Center. Refer to the APIs topic in the Programming category of the Information Center. An application opens a file, operates on the file, and finally closes the file.

When an application changes file data or attributes, the optical file system stores these changes in a temporary system object in OS/400 storage. The optical file system does not update the optical disk until the application closes the file. When two or more applications concurrently change file data or attributes, the optical file system updates the optical disk when the last updating application closes the file. The application may force file and attribute data to optical disk by issuing either the HFS API Force Buffered Data or UNIX-Type (fsync) functions.

The following benefits are realized with this implementation:

- Simulation of read and write access to optical files
- File locking and sharing
- Byte locking and sharing
- Random processing of optical file data
- Variable-length data buffers can be written to the optical file
- Reduction of input and output to the optical disk

Held Optical Files

If the optical file system is unable to update the optical disk during a close function, the operation fails and the file is marked as **held**. The optical file system might still consider the file to be open. If it does, it allows any application that already has the file open to continue operating on it. In any case, no new application can open a file while it remains held.

If the system can correct the condition that caused the failure, and the file is still open, the application may attempt to close the file again. If the close function succeeds, the system no longer holds the file. (If a HFS application specified an open type of *normal*, it cannot access the file through the HFS API any longer. IBM provides online information on the open types that concern the Open Stream File command. Refer to the APIs topic in the Programming category of the Information Center..)

Note: The system does not create held files when files fail to close on UDF formatted media.

Recovery for Held Optical Files

If a close operation fails for an open optical file and the file becomes held, the held file can be handled in one or both of the following ways:

- You can attempt a save request
- You can release the file to allow it to be opened again

However, if the cause for the close failure has been corrected, the file can now be closed as usual, without having to save or release it first. In this situation, the file is automatically saved and released, and the held status is lifted.

After releasing a held file, you can close it if the open file handle is still valid.

Before saving or releasing a held optical file, you can view all of its open instances by selecting option 8 (Display Usage information) from the Work with Held Files Optical Files (WRKHLDOPTF) display. This can be an important step in determining the appropriate actions for the file. For example, before deciding not to save the latest version of a file, it is useful to know if other applications have been making concurrent updates to the same file. Updates would roll back for all users if the file were only released and no further updates were made before the last updating application closed the file.

Saving a Held Optical File

Saving a held optical file physically writes the data and file attributes to the optical disk. You can choose to save to the original volume, directory, and file name specified at open time, or to a new optical file path.

In some situations, you can save the file at the original storage destination. For example, if the file has been opened with a *normal* open type, the file is now inaccessible through the HFS API, rendering the open file handle no longer valid. However, the condition that caused the file to become held might have been corrected, giving you the ability to save the data by specifying the held file as the destination.

If the application specifies a different file path as the destination, the file must not already exist. If appropriate, you can delete such a file before attempting to save to that volume, directory, and file name.

After a held optical file is saved, it should be released to allow the file to be used by future applications.

Releasing a Held Optical File

You can release a held file only if no locks are currently imposed on the file by other active jobs.

Releasing a held optical file clears the held status and allows new applications to open the file. It also releases the optical file system from its obligation to update the optical disk, unless some application makes further updates to the file. Once the file has been released, it may be closed if the user's process is still active.

If one or more applications continue to change a file after it is released, the optical file system attempts to update the optical disk when the last updating application

closes the file. However, if the cause for the close failure has not been corrected, the file can be expected to become held again.

A held file can be released after a save or without any save operation. If a successful save cannot be achieved, you can release the file simply to acknowledge that the data cannot be written to the disk and that this result is being accepted without taking further action aside from closing the file.

If you do not release the held optical file, it remains held even if the process ends, unless an automatic close is successful in saving the file at that time. For held files, this might only happen if the open type is *permanent*, provided that the cause for the earlier close failure has been resolved.

Implementing Held Optical File Functions

Before deciding whether to save or release a held optical file, you might want to view information that can influence save and release decisions. The Work with Held Optical Files utility provides this means, in addition to the ability to save and release. The save and release functions are also available as optical-specific functions of the Control File System HFS API.

The Work with Held Optical Files utility provides a convenient way to list and manage any held optical files on the system. Use the Work with Held Optical Files (WRKHLDOPF) command to access the Work with Held Optical Files display.

Options on the Work with Held Optical Files Display are selected to display the use (open instances) of files, as well as to save and release held files. By default, using option 9 (Save) on the Work with Held Optical Files display causes the automatic release of a held file after it is saved.

The functions provided by option 9 (Save) and option 6 (Release) are also available as the optical-specific functions, Save Held Optical File and Release Held Optical File, of the Control File System API. See “Control File System Functions to Optical Support” on page 67 for details on coding syntax.

Unlike option 9 (Save), the Save Held Optical File function of the Control File System API does not automatically release a held file after it is saved. Therefore, an explicit release request is needed afterwards.

Disabling Held Optical File Support

OS/400 is shipped with held optical file support enabled. If desired, you may disable it by using the Change Optical Attributes (CHGOPTA) CL command. When held optical file support is disabled, a held file is not created when a file has failed to archive to optical disk. When using this option, it is up to the user application to manage recovery procedures for files that failed to archive. Consider the following scenarios:

Scenario #1

Application opens an optical file for write, and then writes data to the file. When you attempt to close the file, it fails because the optical disk is full.

- **Held File Support Enabled**
- The file is still open, but becomes held. The file closes when the job ends if it never closes successfully before the job ends. The file will remain held until it is ‘released’.
- **Held File Support Disabled**

- The file is still open, but is not held. The file closes when the job ends if it never closes successfully before the job ends. The file will not become held, and all resources (virtual optical file) associated with the held file will be freed up.

Scenario #2

Application opens an optical file for write, and then writes data to the file. The application then issues a Force Buffered Data API to ensure the data is safe on nonvolatile storage. The iSeries server then loses power.

- **Held File Support Enabled**
- After the IPL of the iSeries server completes the file exists as a held optical file. All data that was successfully 'forced' to disk is recoverable. In other words, when you save the held file to optical storage, all data written before the Force Buffered Data request will be saved.
- **Held File Support Disabled**
- After the IPL of the iSeries server completes, the file does not exist as a held optical file. All data written to this file on the previous open instance is lost. The force data request had no effect.

It is important to note that when held optical file support is disabled, forcing data to nonvolatile storage is meaningless. This is because data writes to optical storage after the file closes successfully. Force Buffered Data will force the data to OS/400 disk. And you can use the held optical file to recover the data after a power loss. Held files are the only mechanism to recover data forced to nonvolatile storage after a power loss or other unexpected error. Held file support is needed to recover any data from an open instance that closes unsuccessfully. This effects the following application program interfaces.

- Force Buffered Data HFS API (QHFFRCFSF)
This API will still be allowed when held file support is disabled but it will have no effect.
- Synchronize File Changes IFS API (fsync)
This API will still be allowed when held file support is disabled but it will have no effect.
- Synchronous write-through flag on Open Stream File HFS API (QHFOPNSF)
This value will be allowed, but will be treated as asynchronous write-through flag.

Use the CHGOPTA CL command to enable, disable or determine the current status of held optical file support. Once held optical file support is disabled, it remains disabled for all optical users. You must enable held file support for it to become active again. IBM provided online information that describes the CHGOPTA CL command. Refer to the CL topic in the Programming category of the Information Center.

Recovery Examples

- An application attempts to close a file but fails because the media is corrupted. The file is now held. The user saves the file to a different optical volume and releases the file. The file is no longer held.
- A user opens a file and writes to it. Meanwhile, the optical controller is varied off. The application fails in its attempt to close. The file is now held, but the user can still use the open file handle. The user varies on the optical controller and then closes the file. The file is automatically updated with the close because a

release request was **not** issued. Also, the file is no longer held because the second close attempt was successful in saving the file.

Appendix D. Output File Support Structures

This appendix describes the three possible record formats that are created by the Display Optical (DSPOPT) command when output is directed to either an output file or user space.

Only the fields that are set for a LAN volume (volume type 9) are indicated with an asterisk (*). The *Volume Capacity* and *Volume Space Available* fields have a different meaning for LAN volumes; this is because the *Volume Full Threshold* field is not available. The *Volume Capacity* field contains the current total free space on the volume. Total free space equals the free space available for the user's application plus the space reserved by the volume full threshold.

The *Volume Space Available* field contains the user free space on the volume. The user free space equals the total free space less the amount of space reserved by the volume-full threshold.

Output File Structure for Volume Attributes

Following is the record format for QAMODVA:

* CENTURY	CHAR(1)
* DATE	CHAR(6)
* TIME	CHAR(6)
* VOLUME NAME	CHAR(32)
OPTICAL DEVICE	CHAR(10)
* CSI	CHAR(8)
* CSI LIBRARY	CHAR(10)
* AUTHORIZATION LIST	CHAR(10)
INTERNAL VOLUME ID	CHAR(32)
VOLUME SERIAL NUMBER	PACKED(11,0)
* VOLUME TYPE	PACKED(3,0)
VOLUME CCSID	CHAR(2)
* MEDIA TYPE	PACKED(3,0)
MEDIA FORMAT	PACKED(3,0)
VOLUME FULL THRESHOLD	PACKED(5,0)
VOLUME SEQUENCE NUMBER	PACKED(9,0)
VOLUME CREATION DATE	CHAR(7)
VOLUME CREATION TIME	CHAR(6)
VOLUME DESCRIPTION TEXT	CHAR(50)
VOLUME LAST REFERENCE DATE	CHAR(7)
* OPPOSITE SIDE VOLUME NAME	CHAR(32)
VOLUME BLOCK SIZE	PACKED(9,0)
* VOLUME CAPACITY	PACKED(11,0)
* VOLUME SPACE AVAILABLE	PACKED(11,0)
VOLUME LOCATION	CHAR(1)
VOLUME OFFLINE LOCATION	CHAR(50)
VOLUME ACCESS	CHAR(1)
DOUBLE VOLUME MEDIUM	CHAR(1)
DOUBLE-SIDED MEDIUM	CHAR(1)
IPL-CAPABLE	CHAR(1)
LAST VOLUME OF SET	CHAR(1)
RESERVED	CHAR(23)

When the volume type is backup, the following fields are used:

PRIMARY VOLUME NAME	CHAR(32)
PRIMARY VOLUME SERIAL #	PACKED(11,0)
CMPLT RANGE START DATE	CHAR(7)
CMPLT RANGE START TIME	CHAR(6)

CMPLT RANGE END DATE	CHAR(7)
CMPLT RANGE END TIME	CHAR(6)
VOLUME CHANGED END DATE	CHAR(7)
VOLUME CHANGED END TIME	CHAR(6)

When the volume media is CD-ROM, the following fields are applicable:

MODIFICATION DATE	CHAR(7)
MODIFICATION TIME	CHAR(6)
EXPIRATION DATE	CHAR(7)
EXPIRATION TIME	CHAR(6)
EFFECTIVE DATE	CHAR(7)
EFFECTIVE TIME	CHAR(6)
COPYRIGHT INFORMATION	CHAR(37)
ABSTRACT INFORMATION	CHAR(37)
BIBLIOGRAPHIC INFO	CHAR(37)
PUBLISHER KEY	CHAR(1)
PUBLISHER	CHAR(128)
PREPARER KEY	CHAR(1)
PREPARER	CHAR(128)
DATA SPECIFICATION KEY	CHAR(1)
DATA SPECIFICATION	CHAR(128)

The following constants are used in the status fields:

VOLUME LOCATION	
OFFLINE	CHAR(1) CONSTANT("0")
SLOT	CHAR(1) CONSTANT("1")
DRIVE	CHAR(1) CONSTANT("2")
MOVING	CHAR(1) CONSTANT("3")
REMOVED	CHAR(1) CONSTANT("4")
VOLUME ACCESS	
READ ONLY	CHAR(1) CONSTANT("1")
WRITE PROTECTED	CHAR(1) CONSTANT("2")
WRITABLE	CHAR(1) CONSTANT("3")
DOUBLE VOLUME MEDIUM	
NO	CHAR(1) CONSTANT("0")
YES	CHAR(1) CONSTANT("1")
DOUBLE-SIDED MEDIUM	
NO	CHAR(1) CONSTANT("0")
YES	CHAR(1) CONSTANT("1")
IPL-CAPABLE	
NO	CHAR(1) CONSTANT("0")
YES	CHAR(1) CONSTANT("1")
LAST VOLUME OF SET	
NO	CHAR(1) CONSTANT("0")
YES	CHAR(1) CONSTANT("1")
KEY (PUBLISHER, PREPARER, DATA SPECIFICATION)	
CONTAINS DATA	CHAR(1) CONSTANT("0")
CONTAINS FILE NAME	CHAR(1) CONSTANT("1")
VOLUME TYPE	
PRIMARY	PACKED(3,0) CONSTANT(000.)
BACKUP	PACKED(3,0) CONSTANT(001.)
JOURNAL	PACKED(3,0) CONSTANT(002.)
MIRROR	PACKED(3,0) CONSTANT(003.)
UNFORMATTED	PACKED(3,0) CONSTANT(004.)
UNKNOWN	PACKED(3,0) CONSTANT(005.)
SERVER VOLUME	PACKED(3,0) CONSTANT(009.)
MEDIA TYPE	
WORM	PACKED(3,0) CONSTANT(000.)

ERASABLE	PACKED(3,0)	CONSTANT(001.)
CD-ROM	PACKED(3,0)	CONSTANT(002.)
DVD-ROM	PACKED(3,0)	CONSTANT(003.)
DVD-RAM	PACKED(3,0)	CONSTANT(004.)
UNKNOWN	PACKED(3,0)	CONSTANT(009.)
MEDIA FORMAT		
UNINITIALIZED	PACKED(3,0)	CONSTANT(000.)
HPOFS	PACKED(3,0)	CONSTANT(001.)
ISO9660	PACKED(3,0)	CONSTANT(002.)
UNKNOWN	PACKED(3,0)	CONSTANT(003.)
UDF	PACKED(3,0)	CONSTANT(004.)
UDF PARTIAL	PACKED(3,0)	CONSTANT(005.)
CE CARTRIDGE	PACKED(3,0)	CONSTANT(254.)

Output File Structure for Directory Attributes

Following is the record format for QAMODPA:

CENTURY	CHAR(1)
DATE	CHAR(6)
TIME	CHAR(6)
DIRECTORY NAME	CHAR(256)
VOLUME NAME	CHAR(32)
OPTICAL LIBRARY	CHAR(10)
DIR CREATION DATE	CHAR(7)
DIR CREATION TIME	CHAR(6)
RESERVED	CHAR(25)

Output File Structure for File Attributes

Following is the record format for QAMODFA:

CENTURY	CHAR(1)
DATE	CHAR(6)
TIME	CHAR(6)
PATH NAME	CHAR(256)
VOLUME NAME	CHAR(32)
OPTICAL DEVICE	CHAR(10)
FILE SIZE	PACKED(9,0)
FILE CREATION DATE	CHAR(7)
FILE CREATION TIME	CHAR(6)
FILE MODIFICATION DATE	CHAR(7)
FILE MODIFICATION TIME	CHAR(6)
FILE EXPIRATION DATE	CHAR(7)
FILE EXPIRATION TIME	CHAR(6)
CONT FROM PREVIOUS VOL	CHAR(1)
CONT ON NEXT VOLUME	CHAR(1)
STARTING VOLUME ID	CHAR(32)
ATTRIBUTE NAME	CHAR(25)
ATTRIBUTE DATA	CHAR(75)
FILE SIZE2	PACKED(15,0)
RESERVED	CHAR(17)

Note: If the file size is 999,999,999 bytes or less, FILE SIZE and FILE SIZE 2 will both contain the correct size of the file. If the file size is larger than 999,999,999 bytes, FILE SIZE will be set to 999,999,999 and FILE SIZE 2 will contain the correct file size.

Note: If a file has extended file attributes, there will be one record per extended attribute until all attributes of the file have been listed.

Constants used in the status fields.

CONTINUATION INDICATOR
NO CHAR(1)
YES CHAR(1)

Appendix E. Handling Problems That Are Related to Optical Support

This appendix describes how users can expedite solutions to optical related problems.

Common Optical Support Problems and Questions

Following are hints to help with some common problems and questions encountered with optical devices. The data is presented in a question-answer format.

When writing objects, I receive a message indicating that there is not enough optical media space available. However, the volume is not full. What's wrong?

Either the threshold is incorrectly set, the object being stored is bigger than the available space, or the *spares area* is full. Display the volume attributes of the volume you are writing to using the DSPOPT command. Verify that the threshold and space available values are valid. Also, verify that access to the volume access is *writable*, and not *read only*. If it is read only, then the spares area may be full. The spares area is a set of sectors to which data is written when the original sector is damaged.

My backup volume is filling up before all objects from the primary have been stored. What is using up the extra space?

Several situations could cause this to happen:

- Device errors may have occurred when only part of a file was written. When backup was restarted, the complete file was rewritten.
- If the backup volume type is WORM, it might have been initialized multiple times before the backup, thus wasting some volume space.
- If you are performing an incremental backup, you might have selected the wrong option on the SLTFILE parameter (*ALL instead of *CHANGED).
- If you created the primary volume on a pre-Version 2 Release 3 Modification 0 system, and the primary volume is over 98% full, then the Duplicate Optical (DUPOPT) command may be the only choice to backup this volume.
- The primary volume is a 2X media and the backup volume is 1X.

While backing up a platter, the task ends abnormally. When I restart the backup, I receive the OPT1210 message indicating that the directory already exists. However, the directory is not listed when I use the Work with Optical Directories (WRKOPTDIR) command. How can this be?

When the task ended abnormally, the directory was created on the volume, but the internal optical index files had not been updated yet. Remove the backup volume using the RMVOPTCTG command and add it back in using the ADDOPTCTG command with DIR(*YES)). The internal optical index will be updated with the new path.

I received an OPT1115 message indicating that the file is not found when trying to retrieve an object. When I use the Work with Optical Files (WRKOPTF) command, the object is displayed. Why am I unable to retrieve the object?

The optical media may be dirty. Contact your next level of support (hardware) to get the media cleaned.

My application appears to be storing objects correctly, but when I use the Work with Optical Files (WRKOPTF) command, not all of the objects are showing up. Where are the objects going?

The files may be held optical files. Refer to “Held Optical Files” on page 85 for more information on held optical files. In this case, the volume may have reached its threshold. Verify that your application is handling the OPT1345 message (Threshold reached on optical volume) or CPF1F61 message (No space available on media) correctly.

My application program fails with a CPF1F83 message indicating that the file system name /QOPT was not found when I attempt to copy a stream file using the CPYSF command. What is wrong?

The file system portion of the path (/QOPT) must be specified in uppercase characters. The rest of the path can be in either uppercase or lowercase characters.

There is a volume I want to use, but I can't seem to access it. What can I do?

Duplicate volume names might cause this. If the volume is in a LAN server, it may have the same name as a volume in a directly-attached library or another server. If there are duplicate names, only the first volume found is usable.

Messages are occurring which indicate that I should run Reclaim Optical (RCLOPT). From past experience I know a RCLOPT of type *RESET can take a long time. Is there a quicker way to recover?

Yes, there is! First read Appendix A, “Reclaiming the Optical Index Database,” on page 123 to gain a better understanding of the RCLOPT process. Then attempt one of the following.

- Run RCLOPT MLB(device name) OPTION(*SYNC).
- Work with Optical Volumes (WRKOPTVOL) and press F14 (Show extended information). If any volumes show *moving* as the location, then do the following:
 1. Run RCLOPT MLB(device name) OPTION(*UPDATE) VOL(moving volume name).
 2. Refresh the Work with Optical Volumes screen. If any volumes still show up as *moving*, repeat step 1.
- Run RCLOPT MLB(device name) OPTION(*RESET) DIR(*NO)

Note: This choice takes longer than the first two, but by specifying DIR(*NO), it could cut the RCLOPT *RESET time in half.

What is the difference between volumes marked *OFFLINE and those marked *REMOVED?

*OFFLINE entries are volumes in optical devices that are either powered off, varied off, or no-longer connected. *REMOVED entries are volumes that were removed from the optical media library with VOLOPT *KEEP specified.

When I add full optical volumes into my directly-attached optical library by using the ADDOPTCTG command, it takes a long time. Any suggestions?

When volumes are removed using the RMVOPTCTG command, remove them with *KEEP specified in the VOLOPT parameter. The internal optical indexes save all information about these volumes, including the optical directory information. When volumes are added using the ADDOPTCTG command, specify *NO in the DIR parameter. The volumes are added and the directory index is not rebuilt. This speeds up the import process.

Note: This process should not be followed if modifications were made to the removed volumes since the volumes were last removed from this system.

I entered a CD-ROM volume into my CD-ROM device, but I received a volume not found message when I attempted to access it. I did not see any error messages. What went wrong?

Refer to Chapter 2, "CD-ROM and DVD on iSeries Servers," on page 9 for information on loading CD-ROM media. In this case, you probably attempted to access the CD-ROM before it was fully loaded (wait 10-20 seconds after the tray slides in), or an error occurred during the load. Refer to the message queue for QSYSOPR to see if the CD-ROM volume entered successfully.

I can't seem to use my optical LAN servers. What's wrong?

Check that LAN support is active using the DSPOPTSVR command. It may be that there are no LAN devices defined for the system. The ADDOPTSVR command does not run successfully if there are no devices configured and varied on.

I can use some of my LAN servers, but not all of them. What happened?

The ADDOPTSVR command has a Communications Side Information (CSI) parameter. All CSIs which are going to be used must be specified for this parameter. If some CSIs have not been specified, they are not usable until the ADDOPTSVR command is issued for them.

How do I determine if a job that is using an optical LAN volume is actually sending or receiving information from the optical LAN server?

Use the Work with Job (WRKJOB) command to display information about your job. Select option 17 to display the jobs communications status. The display shows the input and output counts for all communications that are used by that job. APPC-CPIC is the communications method that is used with optical LAN servers.

I have issued the Add Optical Server (ADDOPTSVR) command but it never completes. How do I determine what could be wrong?

Start by displaying the messages in the QSYSOPR message queue. If there is some type of communications problem, it is most likely listed here. Next, check the job log for any diagnostic or escape messages. Also, check the status of the APPC controller that defines the optical LAN server. It may be in a failed state or it could be varied off.

Handling Installation Problems of LAN Optical Media Libraries

A majority of the problems you encounter when working with a LAN- attached optical library are experienced either at installation time or are problems with the configuration. Follow the steps in the Table 11 to help debug installation and configuration problems.

Table 11. Finding and Correcting Install Problems of LAN Optical Libraries Checklist

#	Check and Test	Problem Description and Action	Action or Solution
1	Get the CSI name	Do you know the CSI name?	<p>Yes: Go to #2.</p> <p>No/Unsure: Use WRKCSI to view all CSIs. If no CSIs are found, either you need to add the library containing the CSIs, or you need to configure your OS/400 communications.</p>
2	Is LAN Software operational?	From the command line (CALL QCMD) enter ADDOPTSVR, type the name(s) of the CSI, and press enter.	<p>The message ADDOPTSVR completed successfully appears: Go to #3.</p> <p>The message ADDOPTSVR completed successfully appears along with the message Target program not available. Not allowed to try again. Go to #12.</p> <p>The message ADDOPTSVR did not complete successfully appears: Go to #4.</p> <p>The message CPF9801, Object XXXXXX in library *LIBL not found appears: Go to #5.</p> <p>The machine appears to hang: Go to #6.</p>
3	Confirm ADDOPTSVR Operation	Enter WRKOPTVOL CSI(*ALL) and see if volumes are present.	<p>Volumes are Present: LAN appears to be working correctly.</p> <p>No Volumes Present from one or all servers: The OS/400 software is installed correctly. The problem lies in either the OS/400 dataserver communications or the PC-controller software installation or configuration. Check for messages in the QSYSOPR message queue on the OS/400 to further isolate the problem.</p>
4	Investigate why the ADDOPTSVR command failed	Press F10 to include detailed messages.	<p>Controller/Device Related Error: Check the OS/400 communications setup (#7).</p> <p>Other: Contact your next level of support.</p>
5	CSI missing	The CSI selected in the ADDOPTSVR command was not found. Use the WRKCSI command to display a list of available CSIs. Make sure the CSI library is in the library list.	Correct CSI and retry #2.
6	Job hangs	Display the system operator messages.	If message Controller XXXXXX not replying. Remote system or configuration problem appears, enter C to cancel the message. Either the server is down, or communications has not been properly configured. Go to #7.

Table 11. Finding and Correcting Install Problems of LAN Optical Libraries Checklist (continued)

#	Check and Test	Problem Description and Action	Action or Solution
7	PC Communications started?	On the PC controller, select the CM/2 icon and then select the Subsystem Management icon. Press F5 to refresh the screen. Verify that the APPC Attach Manager, Communications Manager Kernel, and SNA Subsystem all indicate started.	Started: Go to #8. Not started: Select "Start communications" in the CM/2 window and verify that all three items are started (this might take a minute). If the items are not started, there is a problem with the PC. Recheck the setup of the PC.
8	Perform Simple Communication	Enter WRKCFGSTS *CTL. Find the controller and device description and vary them both off. Perform a vary-on reset operation for both.	Controller & Device become active: Go to #11. Controller & Device become Vary-on Pending: Go to #9. Controller or device will not vary-on. There is a problem with the OS/400 device, controller, or line description. Recheck your OS/400 setup.
9	Check Line Status	Enter WRKCFGSTS *LIN. Verify that the line is in a vary on pending or active state.	Line vary on pending or active: Go to #10. Line varied off: Vary on the line and go back to #8.
10	Double-Check PC Information	Display the file C:\IBMCOM\LANTRAN.LOG on the PC. Check for error messages, verify LAN speed, and that the address matches the LAN remote adapter address for the controller (using the DSPCTLD command on OS/400).	Information is correct: The controller and device should become active. There is a problem with the OS/400 controller, device, line description, or the connection itself. Recheck the setup. Information is incorrect: Correct the problem and retry communications.
11	Check Volumes Present	Verify that volumes are present using the dynamic console's display volume information on the server, or on a 3431, by doing a directory of the disk.	Volumes Are Present: Go to #2. No Volumes Present: A volume is necessary to perform the rest of this test. Add a volume before continuing. Go to #2.
12	Check Transaction Program	On the PC-controller, select the CM/2 icon, then select the Subsystem Management icon, and then double click on the Display active Configuration line. Select the display menu, then the general SNA sub-menu, followed by the Transaction Definition menu. Confirm that there is a definition for the HFSSRV program and that the parameters are the drive letter for the server, followed by a D for a 3431 device or a L for anything else.	Configuration Correct: Go to #13. Configuration Not Correct: The transaction program has not been correctly defined in the communications manager. The problem lies in the transaction program definition.
13	Check OS/400 Transaction Program Definition	Display the CSI information using the WRKCSI command. Confirm that the transaction program selected is HFSSRV.	Transaction Program Correct: Verify information from #14. The problem is that somehow the transaction program cannot be reached. Otherwise, the system appears to be working correctly. Transaction Program Not Correct: Correct the transaction program and retry this step.

Follow the steps in Table 12 to help debug configuration or installation software problems encountered on the controller for LAN-attached libraries.

PC-Controller Software Installation/Configuration Problems

Table 12. PC-Controller Software Installation/Configuration Problems Checklist

#	Check or Test	Problem Description and Action	Action or Solution
1	Does the system start?	Does the PC start into OS/2?	<p>Yes: Go to #3.</p> <p>Stops After Memory Test and Displays Two Numbers: Restart with the installation diskette and run the auto-configuration program. After the auto-configuration program has completed, restart again and go into the set configuration menu. Verify that the token ring speed is set correctly. Go into the set SCSI menu and set the presence-error reporting to disabled. If the machine still will not start contact your next level of support.</p> <p>Starts to Come Up, but has Problems Loading Device Drivers or OS/2: Go to #2.</p>
2	How much memory does the PC have?	Does the PC have 12 Meg or more of memory?	<p>Yes: The machine has a problem outside the scope of this document. Contact your next level of support.</p> <p>NO: 12 Meg of memory is needed to properly run the dataserer and communications programs.</p>
3	Is the software installed?	Is the HPOFS and HFSSRV software installed?	<p>Yes: Go to #4.</p> <p>No: Install the software diskette that came with the dataserer and follow the included instructions.</p> <p>Unsure: Verify that the PC-controller software is installed.</p>
4	Type of library	Is the controller connected to a 3431 stand-alone, an A22 or A23, or a 022 or 023?	<p>3431 type stand-alone: Go to #5.</p> <p>A22 or A23: Go to #8.</p> <p>022 or 023: Go to #8.</p>
5	3431 Debug	Insert a formatted cartridge in the drive, open an OS/2 window and type DIR D:\. Does the drive light come on and the directory get displayed?	<p>Yes: Go to #6.</p> <p>No: There is a problem with HPOFS. Re-install HPOFS according to the installation instructions.</p>
6	Check HFSSRV Setup	On the PC-controller, select the CM/2 icon, then select the Subsystem Management icon, and then double click on the Display active Configuration line. Select the display menu, then the general SNA sub-menu, followed by the Transaction Definition menu. Confirm that there is a definition for the HFSSRV program and that the parameters are the drive letter for the server, followed by a D.	<p>Configuration Correct: Record the HFSSRV path and go to #7.</p> <p>No Definition Found: CM/2 has not been started or been configured for the transaction program.</p> <p>Configuration Not Correct: The transaction program has not been correctly defined in the communications manager. The problem lies in the transaction program definition.</p>

Table 12. PC-Controller Software Installation/Configuration Problems Checklist (continued)

#	Check or Test	Problem Description and Action	Action or Solution
7	Check HFSSRV Operation	Open an OS/2 window and type the path recorded in the previous step. Does the version number display?	<p>Yes: The 3431 software appears to be installed correctly. If the problems persist, verify the OS/400 and PC communications, or contact your next level of support.</p> <p>No: HFSSRV is incorrectly installed.</p>
8	A22 - A23 - 022 - 023 Debug	On the PC-controller, is the dynamic console window open? If so, select the dynamic console program. Otherwise, click both mouse buttons simultaneously and see if the dynamic console program appears in the window list. If so, select the dynamic console program. Otherwise, see if there is an IBM Optical Library Dataserver icon. If so, select that.	<p>Dynamic Console Program Started: Go to #9.</p> <p>Dynamic Console Program Not Started: The dataserver software has not been correctly installed. Re-install the software. If the problem persists, contact your next level of support.</p>
9	Dynamic console verify	Does the dynamic console window show a white box under the auto-changer and drive icons?	<p>Yes: The PC-controller software appears to be installed correctly. If problems persist, refer to the corresponding service manual or contact your next level of support.</p> <p>No: The dataserver is not functioning. Check the power on the dataserver and the connecting cables. Refer to the corresponding maintenance manual for further debugging, or contact your next level of support.</p>

Collecting Information

Once you have determined that the next level of support must be called, having the following information ready will help speed up the problem analysis process.

- Detailed description of problem, including:
 1. Applications that are running.
 2. Is this a new installation or has the system or application been running?
 3. Is the problem reproducible?
- Type and model number of the dataservers
- Current PTF level
- How many dataservers
- How many volumes
- If LAN, then:
 1. Token ring or Ethernet
 2. LAN speed
 3. How many iSeries servers are attached

Other System Commands

Other system commands to be familiar with to gather pertinent information include:

- The Display Job Log (DSPJOBLOG) command shows commands and related messages for a job while it is still active and has not yet been written.

- The Display Log (DSPLOG) command shows the system history log (QHST). The history log contains information about the operation of the system and the system status.
- The Trace Job (TRCJOB) command controls traces of program calls and returns that occur in the current program, or the job being serviced.
- The Start Service Job (STRSRVJOB) command starts the remote service operation for a specified job so that other service commands can be entered to service the specified job.
- The End Service Job (ENDSRVJOB) command ends the remote job service operation. This command stops the service operation that began when the Start Service Job (STRSRVJOB) command was entered.
- The Analyze Problem (ANZPRB) command allows you to analyze, create problem records for, or report problems that were not detected by the system. If the problem is valid, a fix can be supplied by matching the problem description to an already known problem for which a PTF exists, or an APAR can be created.

IBM provides online information on each of the above commands. Refer to the *CL and APIs* section that is found in the *Programming* category of the **Information Center** at the following Web site - <http://publib.boulder.ibm.com/pubs/html/as400/infocenter.htm>.

Appendix F. Programming Examples

This appendix provides examples of using various interfaces to optical support. Although the examples are basic, they demonstrate the structures and steps necessary to use various APIs in different environments.

The examples demonstrate the following interfaces and environments:

- HFS programming examples using Integrated Language Environment® (ILE) RPG for OS/400
- Integrated file system programming example using Integrated Language Environment (ILE) C for OS/400
- Optical tools using OS/400 commands and CL programs

HFS API Program Examples Using ILE RPG for OS/400

This topic demonstrates how the HFS API can be used with the ILE RPG for OS/400 programming language.

The programming examples demonstrate the following functions:

- Retrieving a path name from an array
- Calling the HFS API to open a stream file
- Calling the HFS API to write a 256-byte buffer passed to the program as a parameter
- Calling the HFS API to close the stream file

IBM provides online information for HFS APIs. Refer to the APIs topic in the Programming category of the Information Center.

Getting a Path and Calling Subroutines

```
E                AR      1   5 36

C      *ENTRY      PLIST
* 2 PARAMETERS - A DATA BUFFER ID AND AN INDEX TO THE ARRAY
C                PARM      DATAIN   256
C                PARM      IDX       10
* MOVE THE ARRAY ELEMENT TO A FIELD CALLED "PATH"
C                MOVE AR,IDX   PATH
* EXECUTE SUBROUTINES TO OPEN, WRITE AND CLOSE A FILE
C                EXSR OPNSF
C      RTCD       IFEQ 0
C                EXSR WRTSF
C                EXSR CLOSF
C                END
C                SETON                      LR

* TABLE/ARRAY . . . . . :  AR
**
/QOPT/MYVOL1/DIRA/FILE
/QOPT/MYVOL1/DIRA/SUBDIRB/FILE
/QOPT/MYVOL1/DIRA/SUBDIRB/C/FILE
/QOPT/MYVOL1/DIRA/SUBDIRB/C/D/FILE
/QOPT/MYVOL1/DIRA/SUBDIRB/C/D/E/FILE
```

Defining Data Structures for Opening Files

```

* PATH LENGTH PARAMETER
IPATHLN      DS
I                                B  1    40PATHL
* OPEN INFORMATION PARAMETER
IOPNINF      DS
I                                1  1  EXISTS
I                                2  2  NOTTHR
I                                3  3  SYNASY
I                                4  4  RSV1
I                                5  5  SHAREM
I                                6  6  ACCESS
I                                7  7  OTYPE
I                                8 10  RSV3
* ATTRIBUTE LENGTH PARAMETER
IATTRLN      DS
I                                B  1    40ATTRL
* RETURN CODE PARAMETER
IRETCD       DS
I                                B  1    40RCLEN
I                                B  5    80RTCD
I                                9   15  CONDTN
I                                16   16  RSV
I                                17  272  MSG
* BYTES TO READ/WRITE PARAMETER
IBYTRDW      DS
I                                B  1    40BYT2RW
* BYTES ACTUALLY READ/WITTEN PARAMETER
IBYTACT      DS
I                                B  1    40BYTARW

```

Opening an Optical File

```

* PARAMETER LIST FOR QHFOPNSF CALL
C          POPNSF  PLIST
C          PARM    FHDLE  16
C          PARM    PATH   36
C          PARM    PATHL
C          PARM    OPNINF
C          PARM    ATRTBL  1
C          PARM    ATTRLN
C          PARM    ACTION  1
C          PARM    RETCD

C* OPEN FILE SUBROUTINE
C          OPNSF   BEGSR
C* FILL IN THE PATH AND ATTRIBUTE LENGTHS
C          Z-ADD36  PATHL          SET PATH LEN=36
C          Z-ADD*ZEROS  ATTRL      ZERO ATTRIBUTE LNGTH
C* FILL IN THE OPNINF PARAMETER
C          MOVE '0'  EXISTS  1    FAIL IF EXISTS
C          MOVE '1'  NOTTHR  1    CREATE IF NOT THERE
C          MOVE '0'  SYNASY  1    ASYNCHRONOUS
C          MOVE *BLANKS  RSV1  1
C          MOVE '1'  SHAREM  1    DENY NONE
C          MOVE '2'  ACCESS  1    READ/WRITE
C          MOVE '0'  OTYPE  1    NORMAL
C          MOVE *BLANKS  RSV3  3
C* CALL THE API TO OPEN THE STREAM FILE
C          CALL 'QHFOPNSF' POPNSF          50
C          OPNEND  ENDSR

```

Writing a File to an Optical Disk

```

* PARAMETER LIST FOR QHFRDSF OR QHFWRTSF CALL
C          PRWSF   PLIST
C          PARM    FHDLE  16

```

```

C          PARM          DATAIN
C          PARM          BYT2RW
C          PARM          BYTARW
C          PARM          RETCD

C* CALL API TO WRITE TO THE FILE
C          WRTSF        BEGSR
C          Z-ADD256      BYT2RW          SET WRITE LENGTH=256
C          CALL 'QHFWRTSF'PRWSF          50
C          WRTEND        ENDSR

```

Closing an Optical File

```

* PARAMETER LIST FOR QHFCLOSF CALL
C          PCLOSF      PLIST
C          PARM          FHDLE    16
C          PARM          RETCD

C* CALL API TO CLOSE THE FILE
C          CLOSF        BEGSR
C          CALL 'QHFCLOSF'PCLOSF      50
C          CLSEND        ENDSR
C* END OF SAMPLE RPG CALL TO THE HFS API

```

Integrated File System Program Examples Using ILE C for OS/400

This topic demonstrates the use of the integrated file system UNIX-type APIs that pertain to the QOPT physical file system and are used with the ILE C for OS/400 programming language.

The programming examples demonstrate the following functions:

- Retrieving optical directory entries
- Creating an optical file
- Writing a file
- Closing a file
- Opening a file
- Reading a file
- Changing the offset into a file

IBM provided online information for the "UNIX-Type APIs". Refer to the APIs topic in the Programming category of the Information Center.

Integrated File System Function Examples

```

/*****/
/*
/* This program demonstrates the use of various integrated file
/* system functions applied to the QOPT physical file system
/* including:
/* chdir() - change current directory
/* close() - close file
/* closedir() - close directory
/* creat() - create file
/* lseek() - seek file (change file offset)
/* open() - open file
/* opendir() - open directory
/* read() - read file
/* readdir() - read directory entry
/* rewinddir() - rewind directory entries
/* stat() - directory statistics
/* write() - write file

```

```

/*                                                                 */
/*****/
#include <stdio.h>
#include <unistd.h>
#include <sys/types.h>
#include <dirent.h>
#include <sys/stat.h>
#include <fcntl.h>

void main (void)
{
    /*****/
    /* local variables, contents and defines */
    /*****/
    char path[294];           /* optical path */
    DIR *dirP;               /* pointer to the directory */
    int filedes;             /* open file descriptor */
    struct dirent *direntP;  /* directory entry structure */
    struct stat info;        /* dir/file information */
    int volume_number;       /* what it says... */
    int rc = 0;              /* function return codes */
    int kk = 0;              /* local counter */
    char data[] = "The quick red fox jumped over the fence";

    /*****/
    /* Retrieve the list of volumes from the QOPT physical file */
    /* system by opening the QOPT pfs root directory and reading the */
    /* directory entries. */
    /*****/
    memset(path,             /* clear path name */
           0x00,
           sizeof(path));
    strcpy(path,             /* set physical file system */
           "/QOPT");
    rc = stat("/QOPT", &info); /* determine number of files */
    if (rc != 0)
        perror("stat() failed:");

    dirP = opendir(path);   /* open the directory */
    if (dirP == NULL)
        perror("opendir() failed:");

    for (kk = 1; kk <= info.st_nlink; kk++)
    {
        direntP = readdir(dirP);
        if (direntP == NULL)
            perror("readdir() failed:");
        printf("%d) %s\n", kk, direntP->d_name);
    }

    /*****/
    /* Prompt user for the volume they want to work with and make it */
    /* the current directory. */
    /*****/
    printf("\nEnter the number the volume you want to work with:\n");
    scanf("%d", &volume_number);

    rewinddir(dirP);        /* beginning of directory */
    for (kk = 1; kk <= volume_number; kk++)
        direntP = readdir(dirP); /* get requested dir. entry */

    strcat(path, "/");
    strcat(path, direntP->d_name);
    rc = chdir(path);       /* set current working dir. */
    if (rc != 0)
        perror("chdir() failed:");
    if (getcwd(path, sizeof(path)) == NULL)

```

```

perror("getcwd() failed:");
printf("\nThe current working directory is: %s\n", path);

rc = closedir(dirP);          /* close the directory      */
if (rc != 0)
perror("closedir() failed:");

/*****
/* Create and open a file write only. If the file exists it
/* will be truncated. The owner will have read, write, and
/* execute authority to the file.
*****/
strcat(path, "/");
printf("\nEnter a file name:\n");
scanf("%s", &path[strlen(path)]);

filedes = creat(path, S_IRWXU);
if (filedes == -1)
{
perror("creat() failed");
return;
}

rc = write(filedes, data, sizeof(data));
if (rc == -1)
perror("write() failed:");

close(filedes);

/*****
/* Read back the file and print it.
*****/
memset(data, 0x00, sizeof(data));
filedes = open(path, O_RDWR);
if (filedes == -1)
{
perror("open() failed");
return;
}

read(filedes, data, sizeof(data));
if (filedes == -1)
{
perror("read() failed");
return;
}
printf("\nThe data written to file is: %s\n", data);

/*****
/* Change the offset into the file and change part of it. Read
/* the entire file, print it out and close the file.
*****/
lseek(filedes, 4, SEEK_SET);
rc = write(filedes, "slow old ", 9);
if (rc == -1)
{
perror("write() failed");
return;
}
lseek(filedes, 18, SEEK_SET);
rc = write(filedes, "went under ", 11);
if (rc == -1)
{
perror("write() failed");
return;
}

```

```

lseek(filedes, 0, SEEK_SET);
read(filedes, data, sizeof(data));
if (filedes == -1)
{
perror("read() failed");
return;
}
printf("\nThe data now is: %s\n", data);

close(filedes);

printf("Done...\n");
return;
}

```

Optical Tools

This topic contains examples of four CL tools that may be used to move spooled files and database members to and from optical storage:

- Copy stream file
- Copy database to optical
- Copy spooled file to optical
- Copy optical to database

Copy Stream File

Command Source

```

/*****/
/*
/* COMMAND NAME:  CPYSTRF
/*
/* COMMAND TITLE:  Copy stream file
/*
/* COMMAND DESCRIPTION:  Copy stream file between two file systems
/*
/*
/*****/
      CMD  PROMPT('Copy Stream File')

      PARM      KWD(SRCFILE) TYPE(*CHAR) LEN(300) MIN(1) +
                MAX(1) PROMPT('Source file name')      +
                VARY(*YES)

      PARM      KWD(TGTFILE) TYPE(*CHAR) LEN(300) MIN(1) +
                MAX(1) PROMPT('Target file name')      +
                VARY(*YES)

      PARM      KWD(RPLFILE) TYPE(*CHAR) LEN(6) DFT(*NO) +
                SPCVAL((*NO '0  ') (*YES '1  '))      +
                PROMPT('Replace existing file')

```

CL Program Source

```

/*****/
/*
/* PROGRAM:  CPYSTRF (Copy stream file)
/*
/*
/* DESCRIPTION:
/* This is the CL program for sample CL command CPYSTRF.  This
/* program can be used to copy stream files between file
/* systems.  The actual copy is done by making a call to
/* the HFS API program QHFCPYSF (Copy stream file).
/*

```



```

/* */
/* */
/* INPUT PARAMETERS: */
/* - Complete source path */
/* Example: /filesystem/directory1/directoryx/file */
/*           /QDLS/DIRA/DIRB/FILE01 */
/*           - or - */
/*           /filesystem/volume/directory1/directoryx/file */
/*           /QOPT/VOLN01/DIRA/DIRB/FILE01 */
/* - Complete target path */
/* Note: Except for the file the path must already exist. */
/* Example: /filesystem/directory1/directoryx/file */
/*           /QDLS/DIRA/DIRB/FILE01 */
/*           - or - */
/*           /filesystem/volume/directory1/directoryx/file */
/*           /QOPT/VOLN01/DIRA/DIRB/FILE01 */
/* - Replace existing target file */
/* *YES - replace existing file */
/* *NO - do not replace existing file */
/* */
/* LOGIC: */
/* - Separate source file length and value */
/* - Ensure source path is converted to upper case */
/* - Separate target file length and value */
/* - Ensure target path is converted to upper case */
/* - Call copy stream file */
/* */
/* */
/* EXAMPLE: */
/* The example will copy document THISWEEK from folder BILLS */
/* to optical volume YEAR1993. The document will be put into */
/* directory /BILLS/DEC as file WEEK50. */
/* Folders are stored in file system DLS (document library services)*/
/* */
/*           CPYSTRF SRCFILE('/QDLS/BILLS/THISWEEK') */
/*           TGTFILE('/QOPT/YEAR1993/BILLS/DEC/WEEK50') */
/*           RPLFILE(*NO) */
/* */
/* */
/*****/

```

```
PGM PARM(&SRCFILE &TGTFILE &CPYINFO);
```

```

/*****/
/* Input parameters */
/*****/
DCL VAR(&SRCFILE); TYPE(*CHAR) LEN(300)
DCL VAR(&TGTFILE); TYPE(*CHAR) LEN(300)
DCL VAR(&CPYINFO); TYPE(*CHAR) LEN(6)

/*****/
/* Program variables */
/*****/
DCL VAR(&SRCLLEN); TYPE(*CHAR) LEN(4) +
VALUE(X'00000000')
DCL VAR(&TGTLLEN); TYPE(*CHAR) LEN(4) +
VALUE(X'00000000')
DCL VAR(&ERRCODE); TYPE(*CHAR) LEN(4) +
VALUE(X'00000000')
DCL VAR(&COUNT); TYPE(*DEC) LEN(5 0)
DCL VAR(&TBL); TYPE(*CHAR) LEN(10) +
VALUE('QSYSTRNTBL')
DCL VAR(&LIB); TYPE(*CHAR) LEN(10) +
VALUE('QSYS ')

/*****/
/* Monitor for any messages sent to this program */
/*****/

```

```

/*****/
MONMSG      MSGID(CPF0000) EXEC(GOTO CMDLBL(DONE))
MONMSG      MSGID(OPT0000) EXEC(GOTO CMDLBL(DONE))

/*****/
/* The HFS API needs to be passed the file and the file length. */
/* By coding the VARY(*YES) parameter on the command definition */
/* for the source and target file we are passed the length of */
/* entered value as a 2 byte binary field which precedes the */
/* actual value entered. */
/*****/

/*****/
/* Separate source file length and file value. Ensure source */
/* file is upper case. */
/*****/
CHGVAR VAR(%SST(&SRCLN 3 2)) VALUE(%SST(&SRCFIL 1 2))
CHGVAR VAR(%SST(&SRCFIL 1 300)) VALUE(%SST(&SRCFIL 3 298))

CHGVAR VAR(&COUNT); VALUE(%BIN(&SRCLN 3 2))
CALL QDCXLATE (&COUNT      +
              &SRCFIL       +
              &TBL          +
              &LIB)

/*****/
/* Separate target file length and file value. Ensure target */
/* file is upper case. */
/*****/
CHGVAR VAR(%SST(&TGTLN 3 2)) VALUE(%SST(&TGTFIL 1 2))
CHGVAR VAR(%SST(&TGTFIL 1 300)) VALUE(%SST(&TGTFIL 3 298))

CHGVAR VAR(&COUNT); VALUE(%BIN(&TGTLN 3 2))
CALL QDCXLATE (&COUNT      +
              &TGTFIL       +
              &TBL          +
              &LIB)

/*****/
/* Call the copy stream file HFS API to copy the source file to */
/* the target file. */
/*****/
CALL QHFCPYSF (&SRCFIL     +
              &SRCLN       +
              &CPYINFO     +
              &TGTFIL     +
              &TGTLN       +
              &ERRCODE)

SNDPGMMSG MSG('CPYSTRF completed successfully')
RETURN

DONE:
SNDPGMMSG  MSGID(OPT0125) MSGF(QSYS/QCPFMSG)      +
          MSGDTA(CPYSTRF) MSGTYPE(*ESCAPE)

RETURN

ENDPGM

```

Copy Database File to Optical File

CL Command Source

```

/*****/
/*
/* COMMAND NAME:  CPYDBOPT
/*
/*****/

```

```

/* COMMAND TITLE: Copy database to optical */
/* */
/* DESCRIPTION: Copy database file to an optical file */
/* */
/*****/
CPYDBOPT: CMD PROMPT('Copy DB to Optical')

          PARM KWD(FRMFILE) TYPE(QUAL1) MIN(1) +
          PROMPT('From file')

          PARM KWD(FRMMBR) TYPE(*NAME) LEN(10) +
          SPCVAL((*FIRST)) EXPR(*YES) MIN(1) +
          PROMPT('From member')

          PARM KWD(TGTFILE) TYPE(*CHAR) LEN(300) +
          MIN(1) EXPR(*YES) +
          PROMPT('Target file')

QUAL1: QUAL TYPE(*NAME) LEN(10)
        QUAL TYPE(*NAME) LEN(10) DFT(*LIBL) +
        SPCVAL((*LIBL) (*CURLIB)) +
        PROMPT('Library')

```

CL Program Source

```

/*****/
/* */
/* PROGRAM: CPYDBOPT (Copy database to Optical) */
/* */
/* */
/* DESCRIPTION: */
/* This is the CL program for sample CL command CPYDBOPT. This */
/* program can be used to copy a member from a database file to */
/* optical storage. */
/* */
/* */
/* DEPENDENCIES: */
/* - The sample command and program CPYSTRF exists. */
/* - There is an existing folder named OPTICAL.FLR */
/* This folder is used for temporary storage when copying */
/* from database to optical. It is assumed that this folder is */
/* empty and that the user will delete anything which gets */
/* copied into it. */
/* */
/* */
/* INPUT PARAMETERS: */
/* - From file */
/* - From member */
/* - Complete target path */
/* Assumption: - Except for the file the complete path currently */
/* exists. */
/* - File does not currently exist. */
/* Example: /filesystem/volume/directory1/directoryx/file */
/* /QOPT/VOLN01/DIRA/DIRB/FILE01 */
/* */
/* */
/* LOGIC: */
/* - Separate file and library */
/* - Copy file to folder */
/* - Build source file */
/* - Copy file from Document Library Service (DLS) to OPT */
/* */
/* */
/* EXAMPLE: */
/* The example will copy member MYMEMBER in file MYFILE in library */
/* MYLIB to optical storage. It will be stored as file */
/* MYFILE.MYMEMBER in directory /MYLIB on volume VOLN01. */
/* */

```

```

/*          CPYDBOPT FRMFILE(MYLIB/MYFILE)          */
/*          FRMMBR(MYMEMBER)                        */
/*          TGTFILE('/QOPT/VOLN01/MYLIB/MYFILE.MYMEMBER') */
/*          */
/*****/
PGM  PARM(&FROMFILE &FROMMBR &TGTFILE);

/*****/
/* Input parameters */
/*****/
DCL  VAR(&FROMFILE);    TYPE(*CHAR)  LEN(20)
DCL  VAR(&FROMMBR);    TYPE(*CHAR)  LEN(10)
DCL  VAR(&TGTFILE);    TYPE(*CHAR)  LEN(300)

/*****/
/* Program variables */
/*****/
DCL  VAR(&FILE);       TYPE(*CHAR)  LEN(10)
DCL  VAR(&LIB);        TYPE(*CHAR)  LEN(10)
DCL  VAR(&SRCFILE);    TYPE(*CHAR)  LEN(28)          +
VALUE('/QDLS/OPTICAL.FLR/xxxxxxxxxx')

/*****/
/* Monitor for all messages sent to this program */
/*****/
MONMSG  MSGID(CPF0000) EXEC(GOTO CMDLBL(DONE))
MONMSG  MSGID(IWS0000) EXEC(GOTO CMDLBL(DONE))
MONMSG  MSGID(OPT0000) EXEC(GOTO CMDLBL(DONE))

/*****/
/* Separate file and library names then copy the DB file to a */
/* PC folder. */
/*****/
CHGVAR VAR(&FILE); VALUE(%SST(&FROMFILE 1 10))
CHGVAR VAR(&LIB);  VALUE(%SST(&FROMFILE 11 10))

CPYTOPCD FROMFILE(&LIB/&FILE);          +
        TOFLR(OPTICAL.FLR)            +
        FROMMBR(&FROMMBR);            +
        TRNTBL(*NONE)

/*****/
/* Complete the source file path name with the member and copy */
/* the stream file from DLS to optical */
/*****/
CHGVAR VAR(%SST(&SRCFILE 19 10)) VALUE(&FROMMBR);

CPYSTRF SRCFILE(&SRCFILE);          +
        TGTFILE(&TGTFILE);

SNDPGMMSG MSG('CPYDBOPT completed successfully')
RETURN

DONE:
SNDPGMMSG  MSGID(OPT0125) MSGF(QSYS/QCPFMSG) +
        MSGDTA(CPYDBOPT) MSGTYPE(*ESCAPE)

RETURN

ENDPGM

```

Copy Spooled File to Optical

CL Command Source

```
/*
/* *****
/*
/* COMMAND NAME:  CPYSPLFOPT
/*
/*
/* COMMAND TITLE: Copy spooled file to optical
/*
/*
/* DESCRIPTION:   Copy spooled file to an optical file
/*
/*
/* *****
CPYSPLFO:  CMD          PROMPT('Copy Spooled File to Optical')

          PARM          KWD(FRMFILE) TYPE(*NAME) LEN(10)      +
          MIN(1)                               +
          PROMPT('From file')

          PARM          KWD(TGTFILE) TYPE(*CHAR)  LEN(300)    +
          MIN(1)  EXPR(*YES)                               +
          PROMPT('Target file')

          PARM          KWD(JOB) TYPE(Q2)
          DFT(*) SNGVAL(*)
          MIN(0) MAX(1)
          PROMPT('Jobname')

          PARM          KWD(SPLNBR) TYPE(*CHAR) LEN(5)        +
          SPCVAL((*ONLY) (*LAST)) DFT(*ONLY) +
          PROMPT('Spool number')

Q2:       QUAL          TYPE(*NAME) LEN(10)
          MIN(1)
          EXPR(*YES)
          QUAL          TYPE(*NAME) LEN(10)
          EXPR(*YES)
          PROMPT('User')
          QUAL          TYPE(*CHAR) LEN(6)
          RANGE(000000 999999)
          EXPR(*YES) FULL(*YES)
          PROMPT('Number')
```

CL Program Source

```
/*
/* *****
/*
/* PROGRAM:  CPYSPLFOPT (Copy Spooled File to Optical)
/*
/*
/*
/* DESCRIPTION:
/* This is the CL program for sample CL command CPYSPLFOPT. This
/* program can be used to copy a spooled file to optical storage.
/*
/*
/*
/* DEPENDENCIES:
/* - The sample command and program CPYDBOPT exists.
/* - The sample command and program CPYSTRF exists.
/* - There is an existing folder named OPTICAL.FLR
/* This folder is used for temporary storage when copying
/* from spooled files to optical. It is assumed that this folder
/* is empty and that the user will delete anything which gets
/* copied into it.
/* - This CL program uses the CL command CPYSPLF to copy the
/* spooled files to a physical file before copying them to
/* optical. When you use the CPYSPLF command to copy
/* a spooled file to a physical file, certain information can
/* be lost or changed. Before using this command please
```

```

/* refer to the CL Reference Book for the limitations and */
/* restrictions of the CPYSPLF command. */
/* - There is an existing file named LISTINGS in library QUSRSYS. */
/* It is assumed that this file contains no existing members */
/* and that any members that are created will be deleted by the */
/* user. The record length of the file is 133. */
/* */
/* */
/* INPUT PARAMETERS: */
/* - From file */
/* Specify the name of the spooled file to be copied. */
/* - Target file */
/* Assumption: Except for the file the path must already exist. */
/* Example: /filesystem/volume/directory1/directoryx/file */
/* /QOPT/VOLN01/DIRA/DIRB/FILE01 */
/* - Job */
/* Specify the name of the job that created the spooled file */
/* which is to be copied. The possible values are: */
/* * The job that issued this command is the job that */
/* created the spooled file. */
/* - or - */
/* job-name Specify the name of the job that created the */
/* spooled file. */
/* user-name Specify the user name that identifies the user */
/* profile under which the job was run. */
/* job-number Specify the system assigned job number. */
/* - Spool number */
/* If there are multiple files for a job specify the files */
/* spool number. */
/* */
/* */
/* LOGIC: */
/* - Separate job into its three parts: job name, user, job number */
/* - Copy spooled files to database */
/* - Copy database to optical */
/* */
/* */
/* EXAMPLE: */
/* The example will copy spooled file QSYSPRT spool number 2 which */
/* the current process has printed to optical storage. */
/* It will be stored on volume YEAR92 in directory */
/* /DEC/WEEK01/MONDAY as file INVOICES */
/* */
/* CPYSPLFO SPLFILE(QSYSPRT) */
/* TGTFILE('/QOPT/YEAR92/DEC/WEEK01/MONDAY/INVOICES') */
/* SPLNBR(2) */
/* */
/* */
/*****

```

```
PGM PARM(&FROMFILE &TGTFILE &JOB &SPLNBR);
```

```

/*****
/* Input parameters */
/*****
DCL VAR(&FROMFILE); TYPE(*CHAR) LEN(10)
DCL VAR(&TGTFILE); TYPE(*CHAR) LEN(300)
DCL VAR(&JOB); TYPE(*CHAR) LEN(26)
DCL VAR(&SPLNBR); TYPE(*CHAR) LEN(5)

/*****
/* Program variables */
/*****
DCL VAR(&JNAME); TYPE(*CHAR) LEN(10)
DCL VAR(&JUSER); TYPE(*CHAR) LEN(10)
DCL VAR(&JNUM); TYPE(*CHAR) LEN(6)

/*****

```

```

/* Monitor for all messages that can be signalled */
/*****/
MONMSG MSGID(CPF0000) EXEC(GOTO CMDLBL(DONE))
MONMSG MSGID(OPT0000) EXEC(GOTO CMDLBL(DONE))

/*****/
/* Separate each part of the job name and call the copy spool */
/* file command using the current job or the specified name. */
/*****/
CHGVAR VAR(&JNAME); VALUE(%SST(&JOB 1 10))
CHGVAR VAR(&JUSER); VALUE(%SST(&JOB 11 10))
CHGVAR VAR(&JNUM); VALUE(%SST(&JOB 21 6))

IF COND(&JNAME *EQ '*') THEN(DO)
  CPYSPLF FILE(&FROMFILE); +
  TOFILE(QUSRSYS/LISTINGS) +
  TOMBR(&FROMFILE); +
  SPLNBR(&SPLNBR); +
  CTLCHAR(*FCFC)
  ENDDO

ELSE DO
  CPYSPLF FILE(&FROMFILE); +
  TOFILE(QUSRSYS/LISTINGS) +
  TOMBR(&FROMFILE); +
  JOB(&JNUM/&JUSER/&JNAME); +
  SPLNBR(&SPLNBR); +
  CTLCHAR(*FCFC)
  ENDDO

/*****/
/* Copy the database file to optical storage */
/*****/
CPYDBOPT FRMFILE(QUSRSYS/LISTINGS) +
  FRMMBR(&FROMFILE); +
  TGTFILE(&TGTFILE);

SNDPGMMSG MSG('CPYSPLFOPT completed successfully')
RETURN

DONE:
SNDPGMMSG MSGID(OPT0125) MSGF(QSYS/QCPFMSG) +
  MSGDTA(CPYSPLFOPT) MSGTYPE(*ESCAPE)

RETURN

ENDPGM

```

Copy Optical to Database

CL Command Source

```

/*****/
/*
/* COMMAND NAME: CPYOPTDB
/*
/* COMMAND TITLE: Copy optical to database
/*
/* DESCRIPTION: Copy optical file to database file
/*
/*****/
CPYOPTDB: CMD PROMPT('Copy Optical to DB ')
          PARM KWD(SRCFILE) TYPE(*CHAR) LEN(300) +
              MIN(1) EXPR(*YES) +
              PROMPT('Source file')

          PARM KWD(TOFILE) TYPE(QUAL1) MIN(1) +

```

```

                                PROMPT('To file')
                                PARM      KWD(TOMBR) TYPE(*NAME) LEN(10)      +
                                SPCVAL((*FIRST)) EXPR(*YES) MIN(1)          +
                                PROMPT('To member')
QUAL1:  QUAL      TYPE(*NAME) LEN(10)
        QUAL      TYPE(*NAME) LEN(10) DFT(*LIBL)                          +
        SPCVAL((*LIBL) (*CURLIB))                                          +
        PROMPT('Library')

```

CL Program Source

```

/*****/
/*
/* PROGRAM:  CPYOPTDB (Copy Optical to Database)
/*
/*
/* DESCRIPTION:
/* This is the CL program for sample CL command CPYOPTDB. This
/* program can be used to copy a file which is on optical
/* storage to a member of an existing file.
/*
/*
/* DEPENDENCIES:
/* - The sample command and program CPYSTRF exist.
/* - There is an existing folder named OPTICAL.FLR
/* This folder is used for temporary storage when copying
/* from optical to database. It is assumed that this folder is
/* empty and that the user will delete anything which gets
/* copied into it.
/*
/*
/* INPUT PARAMETERS:
/* - Complete source path
/* Example: /filesystem/volume/directory1/directoryx/file
/*          /QOPT/VOLN01/DIRA/DIRB/FILE01
/* - To file
/* Assumptions:
/* - Target library already exists.
/* - Target file already exists and has the same attributes
/* as that which contained the original file.
/* - To member
/*
/*
/* LOGIC:
/* - Build target file
/* - Copy file from OPT to Document Library Services (DLS)
/* - Separate file and library
/* - Copy from folder to database file
/*
/*
/* EXAMPLE:
/* The example will copy file invoices which is in directory
/* DEC on volume YEAR1992. INVOICES was originally a spooled file
/* which had a record length of 133. It will be placed in file
/* LISTINGS which is in library QUSRSYS as member INVOCDEC92.
/*
/*          CPYDBOPT TGTFILE('/QOPT/YEAR1992/DEC/INVOICES')
/*          TOFILE(QUSRSYS/LISTINGS)
/*          TOMBR(INVOCDEC92)
/*
/*****/
PGM  PARM(&SRCFILE &TOFILE &TOMBR);

/*****/
/* Input parameters
*/

```



```

/*****
DCL  VAR(&SRCFILE);    TYPE(*CHAR)  LEN(300)
DCL  VAR(&TOFILE);    TYPE(*CHAR)  LEN(20)
DCL  VAR(&TOMBR);     TYPE(*CHAR)  LEN(10)

/*****
/* Program variables */
/*****
DCL  VAR(&FILE);      TYPE(*CHAR)  LEN(10)
DCL  VAR(&LIB);       TYPE(*CHAR)  LEN(10)
DCL  VAR(&TGTFILE);   TYPE(*CHAR)  LEN(28)          +
                                           VALUE('/QDLS/OPTICAL.FLR/xxxxxxxxxx')

/*****
/* Monitor for all messages signalled */
/*****
MONMSG MSGID(CPF0000) EXEC(GOTO CMDLBL(DONE))
MONMSG MSGID(IWS0000) EXEC(GOTO CMDLBL(DONE))
MONMSG MSGID(OPT0000) EXEC(GOTO CMDLBL(DONE))

/*****
/* Build the target file name and copy the stream file from */
/* optical to DLS */
/*****
CHGVAR VAR(%SST(&TGTFILE 19 10)) VALUE(&TOMBR);

CPYSTRF SRCFILE(&SRCFILE);          +
      TGTFILE(&TGTFILE);

/*****
/* Separate the file and library names. Copy the folder to DB. */
/*****
CHGVAR VAR(&FILE); VALUE(%SST(&TOFILE 1 10))
CHGVAR VAR(&LIB);  VALUE(%SST(&TOFILE 11 10))

CPYFRMPCD FROMFLR(OPTICAL.FLR)      +
          TOFILE(&LIB/&FILE);      +
          FROMDOC(&TOMBR);         +
          TOMBR(&TOMBR);           +
          TRNTBL(*NONE)

SNDPGMMSG MSG('CPYOPTDB completed successfully')
RETURN

DONE:
SNDPGMMSG MSGID(OPT0125) MSGF(QSYS/QCPFMSG) +
          MSGDTA(CPYOPTDB) MSGTYPE(*ESCAPE)
RETURN

ENDPGM

```

Bibliography

The following publications and Information Center topics provide additional information about topics described or referred to in this book.

- The APPC Programming book provides information about developing an application that uses the APPC support.
- *iSeries Handbook*, GA19-5486-23, provides a basic understanding of the AS/400 architecture and operating system.
- The Backup and Recovery topic provides information about the different media available to save and protect system data.
- The Communications Configuration book provides information about configuring communications between systems.
- *IBM 3995 Optical Library Dataserver for model Cn0 with 8x drives: Manual*, SC41-5310, provides IBM service representatives with information on how to install, maintain, and repair the IBM 3995 Optical Support, Model 042.
- *IBM 3995 AS/400 Compact Optical Library Dataserver: Operator's Book*, GA32-0140, provides system operators with information on how to perform basic tasks required by an optical media library dataserver.
- *IBM 3995 AS/400 Compact Optical Library Dataserver: Parts Catalog*, SI32-0024, provides information about the parts that can be ordered for the Compact Optical Library Dataserver, Model 042.
- *IBM 3995 Optical Library Dataserver: Introduction and Planning Book*, GA32-0121, provides introductory information on optical storage, and planning instructions for installation and conversion to optical storage.
- *IBM 3995 Optical Library Dataserver: Introduction and Planning Guide for C-Series Models*, GA32-0350, provides introductory information on optical storage, and planning instructions for installation and conversion to optical storage.
- *IBM 3995 Optical Library Dataserver: Media Requirements*, GA32-0146, describes the minimum requirements for media to be used in the optical library dataserver.
- *IBM 3995 LAN Optical Library Dataserver: Reference*, GA32-0147, provides information regarding the use of the supported optical media library models that are attached through a token-ring LAN.
- *IBM 3995 Optical Library Dataserver Reference for C-Series Models*, GA32-0351, provides information regarding the use of the supported optical media library models that are attached through a token-ring LAN.
- *IBM 3995 Optical Library and IBM 3431 Optical Drive APPC Interface*.
- The Local Device Configuration book provides information about configuring local devices on the AS/400 system. This includes information on how to configure the following:
 - Local work station controllers (including twinaxial controllers)
 - Tape controllers
 - Locally attached devices (including twinaxial devices)
- The iSeries Security Reference book provides information about system security concepts, planning for security, and setting up security on the system. This manual does not describe security for specific licensed programs, languages, and utilities.

This book tells how system security support can be used to protect the system and the data from being used by people who do not have the proper

authorization, protect the data from intentional or unintentional damage or destruction, keep security information up-to-date, and set up security on the system.

- The Basic System Operations topic in the Information Center provides information about the system unit control panel, starting and stopping the system, using tapes and diskettes, working with program temporary fixes, as well as handling problems.
- The Work Management topic in the Information Center provides information about creating and changing the work management environment, working with system values, collecting and using performance data to improve system performance.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

| IBM Director of Licensing
| IBM Corporation
| 500 Columbus Avenue
| Thornwood, NY 10594-1785
| U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

| IBM World Trade Asia Corporation
| Licensing
| 2-31 Roppongi 3-chome, Minato-ku
| Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

| IBM may use or distribute any of the information you supply in any way it
| believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Software Interoperability Coordinator, Department 49XA
3605 Highway 52 N
Rochester, MN 55901
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to

Terms and conditions for downloading and printing publications

Permissions for the use of the publications you have selected for download are granted subject to the following terms and conditions and your indication of acceptance thereof.

Personal Use: You may reproduce these Publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these Publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these Publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these Publications, or reproduce, distribute or display these Publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the Publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations. IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE

All material copyrighted by IBM Corporation.

By downloading or printing a publication from this site, you have indicated your agreement with these terms and conditions.

Code disclaimer information

This document contains programming examples.

IBM grants you a nonexclusive copyright license to use all programming code examples from which you can generate similar function tailored to your own specific needs.

All sample code is provided by IBM for illustrative purposes only. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

All programs contained herein are provided to you "AS IS" without any warranties of any kind. The implied warranties of non-infringement, merchantability and fitness for a particular purpose are expressly disclaimed.

Index

Special characters

- *RESET reclaim type 125
- *SYNC reclaim type 124
- *UPDATE reclaim type 125

A

- Add Optical Cartridge display 41
- Add Optical Server (ADDOPTSVR)
 - command 32, 123
- adding optical cartridges, example 19
- adding optical disk cartridges 41
- adding optical LAN device, example 34
- adding optical volumes 41
- ADDOPTSVR (Add Optical Server)
 - command 123
- advanced program-to-program communications (APPC) 28
- APIs
 - HFS 59
 - using categories concurrently 59
- attribute
 - changing 44
 - definition 71
 - displaying 45
 - printing 47
- authorization list
 - securing LAN volumes 31

B

- backup
 - strategy 89
 - using DUOPT 90
- backup strategy, defining 89
- block size, controlling for performance 109

C

- C/400 programming language
 - examples 157
- cartridges
 - labeling new with uninitialized volumes 16
 - labeling with initialized volumes 16
- categories of APIs for optical support 59
- CD-ROM
 - drive, configuring 9
 - iSeries overview 9
 - volume identifiers 10
- Change Optical Volume Attributes (CHGOPTVOLA) display 44
- changing optical volume attributes 44
- CL tools and program examples 160
- collecting information
 - other system commands 153
- command and device dependency 37

- command, CL
 - Add Optical Server (ADDOPTSVR) 32, 123
 - applicability to particular devices 37
 - Create Communications Side Information (CRTCSI) 30
 - Create Controller Description APPC (CRTCTLAPPC) 30
 - Create Device Description (Media Library) (CRTDEVMLB) 15
 - Create Device Description APPC (CRTDEVAPPC) 30
 - Create Line Description Token Ring (CRTLINTRN) 29
 - Display Optical Server (DSPOPTSVR) 32
 - Duplicate Optical (DUOPT) 90
 - Reclaim Optical (RCLOPT) 123
 - Remove Optical Server (RMVOPTSVR) 32
 - Work with Optical Volumes (WRKOPTVOL) command 16
- communications control setup
 - example 29
- communications side information (CSI)
 - objects, defining 30
- configurations, optical system 2
- configuring
 - CD-ROM device description 9
 - optical LAN support 27
 - optical media libraries 15
- control file system
 - Errors from Control File System (GET) 69
 - functions for directly-attached optical 67
 - get file data 69
 - Release Held Optical File function 68
 - retrieve directory information (LAN-attached) 71
 - retrieve volume information (LAN-attached) 68, 70
 - return volume information or update LAN internal volume list 70
 - Save Held Optical File function 67
 - sector read function 68
- controller (PC) problems 152
- conversation allocation and performance 110
- copy database file to optical file, CL tool 162
- Copy Optical Directories (CPYOPTDIR) display 42
- copy optical file to database file, CL tool 167
- copy option, (COPYTYPE) parameter on Copy Optical display 44
- copy spooled file to optical file, CL tool 165
- copy stream file, CL tool 160

- copying
 - attributes from QDLS to QOPT 74
 - attributes from QOPT to QDLS 74
 - file attributes 74
 - optical volume data 42, 44
 - scenario 44
- Create Communications Side Information (CRTCSI) command 30
- Create Controller Description APPC (CRTCTLAPPC) command 30
- Create Device Description (Media Library) (CRTDEVMLB) 15
- Create Device Description APPC (CRTDEVAPPC) command 30
- Create Line Description Token Ring (CRTLINTRN) command 29

D

- default authorization list 31
- defining backup strategy 89
- deleting *REMOVED volumes from optical index database 50
- dependencies, command and device 37
- device description
 - configuring drives 9
 - working with 15
- differences between directly-attached and LAN-attached devices 33
- differentiating among supported file systems 59
- directly-attached optical libraries
 - definition 2
 - differences from LAN-attached devices 33
 - overview 15
 - performance considerations 107
 - CPYOPT considerations 108
 - DUOPT considerations 108
 - effects of expanding buffer I/O 108
 - effects of file attributes 107
 - using schedule timers 108
- directories, definition 8
- directories, effect on performance 107
- directory attributes
 - output file structure 145
- directory objects 60
- disabling
 - held optical file support 139
- Display Optcial DSPOPT 143
- Display Optical Server (DSPOPTSVR) command 32
- Display Optical Volume Attributes display 45
- Display Optical Volume Attributes display, second display 46
- displaying
 - attributes 45
 - optical volumes 38

- displays
 - Add Optical Cartridge 41
 - Change Optical Volume Attributes (CHGOPTVOLA) 44
 - Copy Optical Directories (CPYOPTDIR) 42
 - copy option (COPYTYPE) parameter 44
 - selecting files to copy, (SLTFILE) parameter 44
 - starting date and time (FROMTIME) parameter 44
 - Display Optical Volume Attributes 45
 - Display Optical Volume Attributes, second display 46
 - Initialize Optical Volume 40
 - Remove Optical Volume 49
 - Work with Optical Volumes 38
 - Work with Optical Volumes display, extended information 39
 - Work with Optical Volumes, view 2 39
- drive contention and effect on performance 107
- DSPOPT 143
- Duplicate Optical (DUPOPT) command 90
- DUPOPT (Duplicate Optical) command 90

E

- error messages associated with optical index database 123
- Ethernet optical LAN support 25
- examples
 - adding an optical LAN device 34
 - adding optical cartridges 19
 - CL tools
 - copy database file to optical file 162
 - copy optical file to database file 167
 - copy spooled file to optical file 165
 - copy stream file 160
 - controlling block size for optical LAN performance 109
 - conversation allocation for optical LAN servers 110
 - copying optical volume data 44
 - hierarchical file system 155
 - initializing optical volumes 20
 - integrated file system 157
 - iSeries communications control setup 29
 - recovery of held optical files 140
 - removing optical disk cartridge 22
 - removing optical LAN device 36
 - working with
 - volumes in a LAN device 35
- expanding buffer I/O
 - applications 84
 - description 72
 - HFS API restrictions 73

- expanding buffer I/O (*continued*)
 - tailoring read requests to increase performance 84
- expanding buffer I/O and performance on directly-attached libraries 108
- extended attribute
 - copying 74
 - OPT.CHGATDTTM Attribute 72
 - QOPT.IOMETH Attribute 72

F

- file attributes
 - output file structure 145
 - performance on directly-attached libraries 107
- file management
 - expanding buffer I/O 84
 - managing media capacity on a per-file basis 83
 - media capacity 82
 - path names 87
 - tips and techniques 82
 - volume threshold 82
 - working with directories 86
 - working with files 87
 - working with volumes 85
- file objects 60
- file size and effect on performance 107
- file, definition 8
- files
 - associated with the optical index database 123
 - effect on performance 107
- finding and correcting problems 147
- force buffered data 85
- fsync() APIs 85
- functions for held optical files 139

G

- generic commands for integrated file system support 75
- GO OPTICAL command 8

H

- handling problems
 - collecting information 153
 - general problem solving 147
 - PC-controller 152
 - system commands you can use 153
 - with installation 150
- held optical files
 - definition 137
 - disabling held file scenarios 139
 - disabling held optical file support 139
 - examples of recovery 140
 - recovering 85, 137, 138
 - releasing 138
 - saving 138
 - using functions 139
- hierarchical file system APIs
 - as implemented by optical support 60

- hierarchical file system APIs (*continued*)
 - control file system (QHCTLFS) functions 67
 - expanding buffer I/O restrictions 73
 - fsync() 85
 - LAN-attached devices 26
 - optical functions 60
 - program examples 155
 - using categories concurrently 59
- High Performance Optical File System 82
- HPOFS
 - High Performance Optical File System 82

I

- ILE C/400 programming language
 - examples 157
- implementation of integrated file system support by optical support 75
- Initialize Optical Volume (INZOPT) display 40
- initializing optical volumes 40
- initializing optical volumes, example 20
- installation, handling problems 150
- integrated file system
 - implementation by optical support 75
 - program examples 157
 - support for optical support by generic commands 75
 - support interface to optical support 75
 - TCP/IP file server support 75
 - UNIX-type APIs 75
- internal library index 124
- introduction to optical support 1

L

- labeling cartridges 15
- LAN
 - communications control setup example 29
 - controller setup 30
 - support 25
- LAN-attached and directly-attached devices, differences 33
- LAN-attached optical libraries
 - definition 3
 - performance considerations
 - block size, controlling 109
 - conversation allocation 110
 - minimizing volume mounts 111
- listing volumes 38
- loading CD-ROM
 - verifying 10

M

- managing media capacity on a per-file basis 83
- media capacity 82

N

nearline, definition 85

O

objects

directory 60
file 60

offline, definition 85

online, definition 85

operating information 9

OPT.CHGATDTM Attribute 72

optical auditing 56

optical directories

considerations 86

optical file system implementation of HFS

APIs 60

optical files

considerations 87

held 85

how they are manipulated 137

how they become held 137

recovering held 137

optical functions provided by HFS

APIs 60

optical index database

associated error messages 123

database files 123

definition 123

description 123

internal library index 124

reclaim types 124, 125

reclaiming 123

reclaiming in a CD-ROM

environment 124

reclaiming on LAN-attached optical

media 123

time required to reclaim 126

optical LAN support

getting started 25

optical media libraries, configuring 15

optical media library, definition 2

optical media types 1

optical server, definition 4

optical storage

directories 8

files 8

using and managing 8

volumes 7

optical support

backup 89

categories of APIs 59

control file system functions 67

file management, tips and

techniques 82

handling installation problems 150

handling problems 147

implementation of HFS APIs 60

implementation of integrated file

system support 75

integrated file system support 75

introduction 1

media capacity 82

questions and answers 147

reclaiming the optical index

database 123

optical support (*continued*)

releasing a held optical file 138

saving held optical files 138

supported hardware 4

unique special attributes 72

volume threshold 82

optical system configurations 2

optical tools 160

optical volumes

adding 41

changing attributes 44

considerations 85

copying data 42

deleting *REMOVED volumes from

optical index database 50

displaying attributes 45

initializing 40

listing 38

printing attributes 47

removing 49

renaming 41

organization of optical storage

directories 8

files 8

volumes 7

output file structure

directory attributes 145

file attributes 145

volume attributes 143

Output File Support Structures 143

P

path name

for volumes in directly-attached

devices 87

for volumes in LAN-attached

devices 88

path, definition 87

PC-controller problems 152

performance

directly-attached libraries

effects of CPYOPT and DUPOPT

commands 108

effects of expanding buffer

I/O 108

effects of file attributes 107

using schedule timers 108

drive contention 107

expanding buffer I/O 84

file size 107

LAN-attached libraries

block size, controlling 109

conversation allocation 110

minimizing volume mounts 111

number of directories and files 107

read requests 84

volume dismounting 107

volume mounting 107

printing

attributes 47

problem handling 147

programming examples 155

Q

QACCDTTM attribute 72

QALCSIZE attribute 71

QAMOPVR optical index database

file 123

QAMOVAR optical index database

file 123

QCRTDTTM attribute 72

QFILATTR attribute 72

QHFCTLFS (control file system)

functions 67

QWRDTTM attribute 72

R

RCLOPT (Reclaim Optical)

command 123

read requests, tailoring for

performance 84

reclaim internal library index reclaim

type 125

Reclaim Optical (RCLOPT)

command 123

reclaim optical directory index

option 125

reclaim types

*RESET 125

*SYNC 124

*UPDATE 125

choosing 126

reclaim internal library index 125

reclaim optical directory index

option 125

synchronize volume index 124

update volume index 125

reclaiming optical index database

CD-ROM only environment 124

description 123

LAN-attached media 123

time required 126

Release Held Optical File function 68

releasing a held optical file 138

Remove Optical Server (RMVOPTSVR)

command 32

Remove Optical Volume display 49

removing optical cartridges 49

removing optical disk cartridge,

example 22

removing optical LAN device,

example 36

removing optical volumes 49

renaming optical volumes 41

RLSHLDOPTF (Release Held Optical File)

function 68

root directory, definition 8

RPG/400 programming language

examples 155

S

Save Held Optical File function 67

SAVHLDOPTF (Save Held Optical File)

function 67

saving held optical files 138

scenarios

copying data 44

- schedule timers, and performance 108
- sector read function 68
- security
 - volumes 31
- selecting files to copy, (SLTFILE)
 - parameter on Copy Optical display 44
- setting write protection 17
- situations in which expanding buffer I/O is useful 84
- special attributes unique to optical file system 72
- standard attribute
 - copying 74
 - definition and restrictions 71
 - QACDDTTM Attribute 72
 - QALCSIZE Attribute 71
 - QCRDTTM Attribute 72
 - QFILATTR Attribute 72
 - QWRDTTM Attribute 72
- starting date and time (FROMTIME)
 - parameter on Copy Optical display 44
- subdirectory, definition 8
- supported hardware for optical storage 4
- Supported Media Format
 - DVD-ROM 9
 - ISO9660 9
 - primary volume descriptor 9
 - PVD 9
 - UDF 9
 - Universal Disk Format 9
- synchronize volume index reclaim type 124

T

- TCP/IP file server support 75
- time required to complete reclaim optical index 126
- token-ring optical LAN support 25
- troubleshooting 147
- types of optical media 1

U

- uninitialized volumes, defined 20
- UNIX-type APIs 75
- unloading CD-ROM
 - verifying 10
- update volume index reclaim type 125
- using and managing optical storage 8
- using API categories concurrently 59

V

- Vary Configuration (VRYCFG)
 - command 9
- varying on CD-ROM
 - verifying 10
- verifying successful loading, unloading, and varying on of CD-ROM 10
- volume attributes
 - output file structure 143
- volume identifiers
 - definition 10

- volume identifiers (*continued*)
 - relationship to loading, unloading, and varying on CD-ROM 10
- volume mounting
 - minimizing for performance on LANs 111
 - performance 107
- volume security 31
- volume threshold 82
- volumes
 - definition 7
 - uninitialized 20

W

- Work with Optical Volumes (WRKOPTVOL) command 16
- Work with Optical Volumes display 38
- Work with Optical Volumes display, extended information 39
- Work with Optical Volumes display, view 2 39
- working with
 - optical volumes 38
 - volumes in a LAN device, example 35
- write protection 17

Readers' Comments — We'd Like to Hear from You

iSeries
Optical Support
Version 5

Publication No. SC41-5310-04

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM CORPORATION
ATTN DEPT 542 IDCLERK
3605 Highway 52 N
ROCHESTER MN 55901-7829



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA

SC41-5310-04

