


**Database Guide**

Version 2







Application System/400

SC41-9659-02

## **Database Guide**

Version 2

**Take Note!**

Before using this information and the product it supports, be sure to read the general information under "Notices" on page xi.

**Third Edition (November 1993)**

This edition applies to the licensed program IBM Operating System/400, (Program 5738-SS1), Version 2 Release 3 Modification 0, and to all subsequent releases and modifications until otherwise indicated in new editions. This major revision makes obsolete SC41-9659-1. Make sure you are using the proper edition for the level of the product.

Order publications through your IBM representative or the IBM branch serving your locality. Publications are not stocked at the address given below.

A Customer Satisfaction Feedback form for readers' comments is provided at the back of this publication. If the form has been removed, you can mail your comments to:

Attn Department 245  
IBM Corporation  
3605 Highway 52 N  
Rochester, MN 55901-7899 USA

or you can fax your comments to:

United States and Canada: 800+937-3430  
Other countries: (+1)+507+253-5192

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you or restricting your use of it.

© **Copyright International Business Machines Corporation 1991, 1993. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.



# Contents

<b>Notices</b> .....	xi
Trademarks and Service Marks .....	xi
<b>About This Guide</b> .....	xiii
Who Should Use This Guide .....	xiii
<b>Summary of Changes</b> .....	xv

---

## Part 1. Setting Up Database Files

<b>Chapter 1. General Considerations</b> .....	1-1
Describing Database Files .....	1-1
Dictionary-Described Data .....	1-2
Methods of Describing Data to the System .....	1-3
Describing a Database File to the System .....	1-4
Describing Database Files Using DDS .....	1-5
Describing the Access Path for the File .....	1-15
Database File Creation: Introduction .....	1-24
Database File and Member Attributes: Introduction .....	1-25
File Name and Member Name (FILE and MBR) Parameters .....	1-25
Physical File Member Control (DTAMBR) Parameter .....	1-25
Source File and Source Member (SRCFILE and SRCMBR) Parameters ..	1-25
Database File Type (FILETYPE) Parameter .....	1-25
Maximum Number of Members Allowed (MAXMBRS) Parameter .....	1-25
Where to Store the Data (UNIT) Parameter .....	1-26
Frequency of Writing Data to Auxiliary Storage (FRCRATIO) Parameter ..	1-26
Frequency of Writing the Access Path (FRCACCPH) Parameter .....	1-26
Check for Record Format Description Changes (LVLCHK) Parameter ..	1-27
Current Access Path Maintenance (MAINT) Parameter .....	1-27
Access Path Recovery (RECOVER) Parameter .....	1-29
File Sharing (SHARE) Parameter .....	1-29
Locked File or Record Wait Time (WAITFILE and WAITRCD) Parameters ..	1-30
Public Authority (AUT) Parameter .....	1-30
System on Which the File Is Created (SYSTEM) Parameter .....	1-30
File and Member Text (TEXT) Parameter .....	1-30
Coded Character Set Identifier (CCSID) Parameter .....	1-30
Sort Sequence (SRTSEQ) Parameter .....	1-30
Language Identifier (LANGID) Parameter .....	1-31
<b>Chapter 2. Setting Up Physical Files</b> .....	2-1
Creating a Physical File .....	2-1
Specifying Physical File and Member Attributes .....	2-1
Expiration Date .....	2-1
Size of the Physical File Member .....	2-2
Storage Allocation .....	2-2
Method of Allocating Storage .....	2-2
Record Length .....	2-3
Deleted Records .....	2-3
Physical File Capabilities .....	2-4

Source Type .....	2-4
<b>Chapter 3. Setting Up Logical Files</b> .....	3-1
Describing Logical File Record Formats .....	3-1
Describing Field Use for Logical Files .....	3-3
Deriving New Fields from Existing Fields .....	3-4
Describing Floating-Point Fields in Logical Files .....	3-6
Describing Access Paths for Logical Files .....	3-7
Selecting and Omitting Records Using Logical Files .....	3-8
Using Existing Access Paths .....	3-13
Creating a Logical File .....	3-15
Creating a Logical File with More Than One Record Format .....	3-16
Logical File Members .....	3-20
Join Logical File Considerations .....	3-23
Basic Concepts of Joining Two Physical Files (Example 1) .....	3-23
Setting Up a Join Logical File .....	3-32
Using More Than One Field to Join Files (Example 2) .....	3-33
Reading Duplicate Records in Secondary Files (Example 3) .....	3-34
Using Join Fields Whose Attributes Are Different (Example 4) .....	3-36
Describing Fields That Never Appear in the Record Format (Example 5) .....	3-37
Specifying Key Fields in Join Logical Files (Example 6) .....	3-39
Specifying Select/Omit Statements in Join Logical Files .....	3-40
Joining Three or More Physical Files (Example 7) .....	3-40
Joining a Physical File to Itself (Example 8) .....	3-42
Using Default Data for Missing Records from Secondary Files (Example 9) .....	3-43
A Complex Join Logical File (Example 10) .....	3-45
Performance Considerations .....	3-46
Data Integrity Considerations .....	3-47
Summary of Rules for Join Logical Files .....	3-47
<b>Chapter 4. Database Security</b> .....	4-1
File and Data Authority .....	4-1
Object Operational Authority .....	4-1
Object Existence Authority .....	4-1
Object Management Authority .....	4-1
Data Authorities .....	4-2
Public Authority .....	4-3
Database File Capabilities .....	4-3
Using Logical Files to Secure Data .....	4-4

---

## Part 2. Processing Database Files in Programs

<b>Chapter 5. Run Time Considerations</b> .....	5-1
File and Member Name .....	5-1
File Processing Options .....	5-2
Specifying the Type of Processing .....	5-2
Specifying the Initial File Position .....	5-3
Reusing Deleted Records .....	5-3
Ignoring the Keyed Sequence Access Path .....	5-4
Delaying End of File Processing .....	5-4
Specifying the Record Length .....	5-4
Ignoring Record Formats .....	5-5
Determining If Duplicate Keys Exist .....	5-5

Data Recovery and Integrity	5-5
Protecting Your File with the Journaling and Commitment Control	5-5
Writing Data and Access Paths to Auxiliary Storage	5-6
Checking Changes to the Record Format Description	5-6
Checking for the File's Expiration Date	5-6
Preventing the Job from Changing Data in the File	5-6
Sharing Database Files Across Jobs	5-6
Record Locks	5-7
File Locks	5-8
Member Locks	5-8
Record Format Data Locks	5-8
Sharing Database Files in the Same Job or Activation Group	5-8
Open Considerations for Files Shared in a Job or Activation Group	5-9
Input/Output Considerations for Files Shared in a Job or Activation Group	5-10
Close Considerations for Files Shared in a Job or Activation Group	5-11
Sequential-Only Processing	5-16
Open Considerations for Sequential-Only Processing	5-16
Input/Output Considerations for Sequential-Only Processing	5-18
Close Considerations for Sequential-Only Processing	5-19
Run Time Summary	5-19
Storage Pool Paging Option Effect on Database Performance	5-21
<b>Chapter 6. Opening a Database File</b>	<b>6-1</b>
Opening a Database File Member	6-1
Using the Open Database File (OPNDBF) Command	6-1
Using the Open Query File (OPNQRYF) Command	6-3
Using an Existing Record Format in the File	6-4
Using a File with a Different Record Format	6-6
OPNQRYF Examples	6-8
The Zero Length Literal and the Contains (*CT) Function	6-9
Selecting Records without Using DDS	6-9
Specifying a Keyed Sequence Access Path without Using DDS	6-21
Specifying Key Fields from Different Files	6-22
Dynamically Joining Database Files without DDS	6-23
Handling Missing Records in Secondary Join Files	6-26
Unique-Key Processing	6-27
Defining Fields Derived from Existing Field Definitions	6-28
Handling Divide by Zero	6-31
Summarizing Data from Database File Records (Grouping)	6-32
Final Total-Only Processing	6-34
Controlling How the System Runs the Open Query File Command	6-35
Considerations for Creating a File and Using the FORMAT Parameter	6-37
Considerations for Arranging Records	6-37
Considerations for DDM Files	6-38
Considerations for Writing a High-Level Language Program	6-38
Messages Sent When the Open Query File (OPNQRYF) Command Is Run	6-38
Using the Open Query File (OPNQRYF) Command for More Than Just	
Input	6-39
Date, Time, and Timestamp Comparisons Using the OPNQRYF Command	6-40
Date, Time, and Timestamp Arithmetic Using OPNQRYF CL Command	6-41
Using the Open Query File (OPNQRYF) Command for Random	
Processing	6-45
Performance Considerations	6-45
Performance Considerations for Sort Sequence Tables	6-48

Performance Comparisons with Other Database Functions . . . . .	6-51
Considerations for Field Use . . . . .	6-52
Considerations for Files Shared in a Job . . . . .	6-53
Considerations for Checking If the Record Format Description Changed . . . . .	6-53
Other Run Time Considerations . . . . .	6-54
Copying from an Open Query File . . . . .	6-54
Typical Errors When Using the Open Query File (OPNQRYF) Command . . . . .	6-55
<b>Chapter 7. Basic Database File Operations . . . . .</b>	<b>7-1</b>
Setting a Position in the File . . . . .	7-1
Reading Database Records . . . . .	7-2
Reading Database Records Using an Arrival Sequence Access Path . . . . .	7-2
Reading Database Records Using a Keyed Sequence Access Path . . . . .	7-3
Waiting for More Records When End of File Is Reached . . . . .	7-5
Releasing Locked Records . . . . .	7-8
Updating Database Records . . . . .	7-8
Adding Database Records . . . . .	7-9
Identifying Which Record Format to Add in a File with Multiple Formats . . . . .	7-9
Using the Force-End-Of-Data Operation . . . . .	7-11
Deleting Database Records . . . . .	7-11
<b>Chapter 8. Closing a Database File . . . . .</b>	<b>8-1</b>
<b>Chapter 9. Handling Database File Errors in a Program . . . . .</b>	<b>9-1</b>

---

### Part 3. Managing Database Files

<b>Chapter 10. Managing Database Members . . . . .</b>	<b>10-1</b>
Member Operations Common to All Database Files . . . . .	10-1
Adding Members to Files . . . . .	10-1
Changing Member Attributes . . . . .	10-2
Renaming Members . . . . .	10-2
Removing Members from Files . . . . .	10-2
Physical File Member Operations . . . . .	10-2
Initializing Data in a Physical File Member . . . . .	10-2
Clearing Data from Physical File Members . . . . .	10-3
Reorganizing Data in Physical File Members . . . . .	10-3
Displaying Records in a Physical File Member . . . . .	10-5
<b>Chapter 11. Changing Database File Descriptions and Attributes . . . . .</b>	<b>11-1</b>
Effect of Changing Fields in a File Description . . . . .	11-1
Changing a Physical File Description and Attributes . . . . .	11-2
Changing a Logical File Description and Attributes . . . . .	11-5
<b>Chapter 12. Using Database Attribute and Cross-Reference Information . . . . .</b>	<b>12-1</b>
Displaying Information about Database Files . . . . .	12-1
Displaying Attributes for a File . . . . .	12-1
Displaying the Descriptions of the Fields in a File . . . . .	12-1
Displaying the Relationships between Files on the System . . . . .	12-2
Displaying the Files Used by Programs . . . . .	12-3
Displaying the System Cross-Reference Files . . . . .	12-4
Writing the Output from a Command Directly to a Database File . . . . .	12-4

<b>Chapter 13. Database Recovery Considerations</b> . . . . .	13-1
Database Save and Restore . . . . .	13-1
Considerations for Save and Restore . . . . .	13-1
Database Data Recovery . . . . .	13-2
Journal Management . . . . .	13-2
Transaction Recovery through Commitment Control . . . . .	13-3
Force-Writing Data to Auxiliary Storage . . . . .	13-4
Access Path Recovery . . . . .	13-5
Saving Access Paths . . . . .	13-5
Restoring Access Paths . . . . .	13-5
Journaling Access Paths . . . . .	13-8
Other Methods to Avoid Rebuilding Access Paths . . . . .	13-9
Database Recovery after an Abnormal System End . . . . .	13-9
Storage Pool Paging Option Effect on Database Recovery . . . . .	13-11
<b>Chapter 14. Using Source Files</b> . . . . .	14-1
Source File Concepts . . . . .	14-1
Creating a Source File . . . . .	14-1
IBM-Supplied Source Files . . . . .	14-2
Source File Attributes . . . . .	14-2
Creating Source Files without DDS . . . . .	14-3
Creating Source Files with DDS . . . . .	14-3
Working with Source Files . . . . .	14-3
Using the Source Entry Utility . . . . .	14-4
Using Device Source Files . . . . .	14-4
Copying Source File Data . . . . .	14-4
Using Source Files in a Program . . . . .	14-6
Creating an Object Using a Source File . . . . .	14-6
Creating an Object from Source Statements in a Batch Job . . . . .	14-7
Determining Which Source File Member Was Used to Create an Object . . . . .	14-8
Managing a Source File . . . . .	14-8
Changing Source File Attributes . . . . .	14-8
Reorganizing Source File Member Data . . . . .	14-9
Determining When a Source Statement Was Changed . . . . .	14-9
Using Source Files for Documentation . . . . .	14-9
<b>Appendix A. Database File Sizes</b> . . . . .	A-1
<b>Appendix B. Double-Byte Character Set (DBCS) Considerations</b> . . . . .	B-1
DBCS Field Data Types . . . . .	B-1
DBCS Constants . . . . .	B-1
DBCS Field Mapping Considerations . . . . .	B-2
DBCS Field Concatenation . . . . .	B-2
DBCS Field Substring Operations . . . . .	B-3
Comparing DBCS Fields in a Logical File . . . . .	B-3
Using DBCS Fields in the Open Query File (OPNQRYF) Command . . . . .	B-4
Using the Wildcard Function with DBCS Fields . . . . .	B-4
Comparing DBCS Fields Through OPNQRYF . . . . .	B-4
Using Concatenation with DBCS Fields through OPNQRYF . . . . .	B-5
Using Sort Sequence with DBCS . . . . .	B-5
<b>Appendix C. Database Lock Considerations</b> . . . . .	C-1
<b>Appendix D. Design Guidelines for OPNQRYF Performance</b> . . . . .	D-1

Overview . . . . .	D-1
Definition of Terms . . . . .	D-1
OS/400 Query Component . . . . .	D-2
Data Management Methods . . . . .	D-4
Access Path . . . . .	D-4
Access Method . . . . .	D-5
The Optimizer . . . . .	D-10
Implementation Cost Estimation . . . . .	D-10
Access Plan and Validation . . . . .	D-12
Optimizer Decision-Making Rules . . . . .	D-12
Join Optimization . . . . .	D-13
Optimizer Messages . . . . .	D-14
Miscellaneous Tips and Techniques . . . . .	D-17
Avoiding Too Many Indexes . . . . .	D-18
Include Selection fields in Ordering Criteria . . . . .	D-18
ORDER BY and ALWCPYDTA . . . . .	D-19
Index Usage with the %WLDCRD Function . . . . .	D-20
Join Optimization . . . . .	D-20
Avoid Numeric Conversion . . . . .	D-22
Avoid String Truncation . . . . .	D-23
Avoid Arithmetic Expressions . . . . .	D-23
<b>Bibliography . . . . .</b>	<b>H-1</b>
<b>Index . . . . .</b>	<b>X-1</b>

# Figures

	1-1.	DDS for a Physical File (ORDHDRP)	1-6
	1-2.	DDS for a Simple Logical File (ORDHDRL)	1-8
	1-3.	DDS for a Field Reference File (DSTREFP)	1-11
	1-4.	DDS for a Physical File (ORDHDRP) Built from a Field Reference File	1-13
	1-5.	DDS for a Logical File (CUSMSTL)	1-14
	1-6.	DDS for a Logical File (CUSTMSTL1) Sharing a Record Format	1-14
	1-7.	MAINT Values	1-28
	1-8.	Recovery Options	1-29
	3-1.	Simple Logical File	3-1
	3-2.	Simple Logical File with Fields Specified	3-1
	3-3.	Three Ways to Code Select/Omit Function	3-9
	3-4.	Physical and Logical Files Before Save and Restore	3-14
	3-5.	Physical and Logical Files After Save and Restore	3-15
	3-6.	DDS for a Physical File (ORDDTLP) Built from a Field Reference File	3-16
	3-7.	DDS for a Physical File (ORDHDRP) Built from a Field Reference File	3-17
	3-8.	DDS for the Logical File ORDFILL	3-17
	3-9.	DDS Example for Joining Two Physical Files	3-25
	3-10.	DDS Example Using the JDUPSEQ Keyword	3-34
	5-1.	Database Processing Options Specified on CL Commands	5-19
	5-2.	Database Processing Options Specified in Programs	5-21
	6-1.	Valid Data Type Comparisons for the OPNQRYF Command	6-11
	6-2.	The STAFF File	6-20
	6-3.	Using the *HEX Sort Sequence	6-20
	6-4.	Using the Shared-Weight Sort Sequence	6-20
	6-5.	Using the Unique-Weight Sort Sequence	6-21
	13-1.	Relationship of Access Path, Maintenance, and Recovery	13-11
	B-1.	Valid Comparisons for DBCS Fields in a Logical File	B-3
	B-2.	Valid Comparisons for DBCS Fields through the OPNQRYF Command	B-4
	C-1.	Database Functions and Locks	C-1
	D-1.	Methods of Accessing AS/400 Data	D-3
	D-2.	Summary of Data Management Methods	D-9





---

## Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. 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 of the intellectual property rights of IBM may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter 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 the IBM Director of Commercial Relations, IBM Corporation, Purchase, NY 10577, U.S.A.

This publication could contain technical inaccuracies or typographical errors.

This publication may refer to products that are announced but not currently available in your country. This publication may also refer to products that have not been announced in your country. IBM makes no commitment to make available any unannounced products referred to herein. The final decision to announce any product is based on IBM's business and technical judgment.

Changes or additions to the text are indicated by a vertical line (|) to the left of the change or addition. Refer to the "Summary of Changes" on page xv for a summary of changes made to the IBM Operating System/400 and how they are described in this publication.

This publication 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.

This publication contains small programs that are furnished by IBM as simple examples to provide an illustration. 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". THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE EXPRESSLY DISCLAIMED.

---

## Trademarks and Service Marks

The following terms, denoted by an asterisk (\*) in this publication, are trademarks of the IBM Corporation in the United States or other countries or both:

AD/Cycle  
Advanced Function Printing  
Advanced Peer-to-Peer Networking  
| AFP  
Application System/400  
APPN  
AS/400  
C/400  
CallPath  
CICS/400  
COBOL/400  
DataHub  
FORTRAN/400  
GDDM  
IBM  
ILE  
Information Warehouse  
Integrated Language Environment

OfficeVision  
OfficeVision/400  
Operating System/400  
Operational Assistant  
OS/2  
OS/400  
PrintManager  
PS  
Q & A  
RM/COBOL-85  
RPG/400  
RUMBA  
SAA  
SQL/400  
Systems Application Architecture  
SystemView  
400

---

## About This Guide

This guide contains information about the AS/400 database management system and describes how to set up and use a database on the AS/400 system.

This guide does not cover in detail all the database-related capabilities on the AS/400 system. Among the topics not fully described are the relationship of database management to:

- Structured Query Language/400 (SQL/400 language)
- Data description specifications (DDS)
- Control language (CL)
- Interactive data definition utility (IDDU)
- Backup and recovery guidelines and utilities

You may need to refer to other IBM manuals for more specific information about a particular topic. The *Publications Guide*, GC41-9678, provides information on all the manuals in the AS/400 library.

For a list of publications related to this guide, see the “Bibliography.”

---

## Who Should Use This Guide

This guide is intended for the system administrator or programmer who creates and manages files and databases on the AS/400 system. In addition, the guide is intended for programmers who use the database in their programs.

Before using this guide, you should be familiar with the introductory material for using the system. You should also understand how to write a high-level language program for the AS/400 system. Use this guide with the high-level language manuals to get additional database information, tips, and techniques.

If you plan to use the SQL/400 language to create and process your database files or OS/400 IDDU to define data on the system, you should also be familiar with those topics.

If you plan to use the database recovery functions on the AS/400 system, you should be familiar with the topics discussed in the *Basic Backup and Recovery Guide*, SC41-0036, and the *Advanced Backup and Recovery Guide*, SC41-8079.



---

## Summary of Changes

|                   **Expanded Sort Sequence and Language Support:** Support is added for sort  
| sequences with the SRTSEQ parameter for both physical and logical files. Also,  
| the LANGID parameter specifies the language to be used with some of the new  
| sort sequence capabilities. Descriptions of the SRTSEQ and LANGID parameters  
| and examples of their use in sorting are in Chapter 1, "General Considerations."

|                   **Implicitly Shared Access Paths:** An expanded explanation with an example of  
| implicitly shared access paths is added in Chapter 3, "Setting Up Logical Files."



---

## Part 1. Setting Up Database Files

The chapters in this part describe in detail how to set up any AS/400\* database file. This includes describing database files and access paths to the system and the different methods that can be used. The ways that your programs use these file descriptions and the differences between using data that is described in a separate file or in the program itself is also discussed.

This part includes a chapter with guidelines for describing and creating logical files. This includes information on describing logical file record formats and different types of field use using data description specifications (DDS). Also information is included on describing access paths using DDS as well as using access paths that already exist in the system. Information on defining logical file members to separate the data into logical groups is also included in this chapter.

A section on join logical files includes considerations for using join logical files, including examples on how to join physical files and the different ways physical files can be joined. Information on performance, integrity, and a summary of rules for join logical files is also included.

There is a chapter on database security in this part which includes information on security functions such as file security, public authority, restricting the ability to change or delete any data in a file, and using logical files to secure data. The different types of authority that can be granted to a user for a database file and the types of authorities you can grant to physical files is also included.





---

## Chapter 1. General Considerations

This chapter discusses things to consider when you set up any AS/400 database file. Later chapters will discuss unique considerations for setting up physical and logical files.

---

### Describing Database Files

Records in database files can be described in two ways:

- Field level description. The fields in the record are described to the system. Some of the things you can describe for each field include: name, length, data type, validity checks, and text description. Database files that are created with field level descriptions are referred to as externally described files.
- Record level description. Only the length of the record in the file is described to the system. The system does *not* know about fields in the file. These database files are referred to as program-described files.

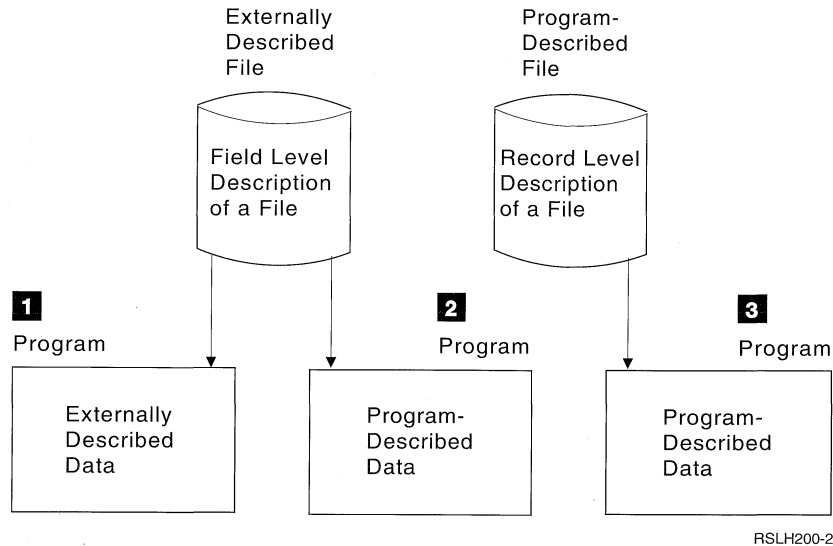
Regardless of whether a file is described to the field or record level, you must describe and create the file before you can compile a program that uses that file. That is, the file must exist on the system before you use it.

Programs can use file descriptions in two ways:

- The program uses the field-level descriptions that are part of the file. Because the field descriptions are external to the program itself, the term, externally described data, is used.
- The program uses fields that are described in the program itself; therefore, the data is called program-described data. Fields in files that are only described to the record level must be described in the program using the file.

Programs can use either externally described or program-described files. However, if you choose to describe a file to the field level, the system can do more for you. For example, when you compile your programs, the system can extract information from an externally described file and automatically include field information in your programs. Therefore, you do not have to code the field information in each program that uses the file.

The following figure shows the typical relationships between files and programs on the AS/400 system:



- 1** The program uses the field level description of a file that is defined to the system. At compilation time, the language compiler copies the external description of the file into the program.
- 2** The program uses a file that is described to the field level to the system, but it does not use the actual field descriptions. At compilation time, the language compiler does not copy the external description of the file into the program. The fields in the file are described in the program. In this case, the field attributes (for example, field length) used in the program must be the same as the field attributes in the external description.
- 3** The program uses a file that is described only to the record level to the system. The fields in the file must be described in the program.

Externally described files can also be described in a program. You might want to use this method for compatibility with previous systems. For example, you want to run programs on the AS/400 system that originally came from a traditional file system. Those programs use program-described data, and the file itself is only described to the record level. At a later time, you describe the file to the field level (externally described file) to use more of the database functions available on the system. Your old programs, containing program-described data, can continue to use the externally described file while new programs use the field-level descriptions that are part of the file. Over time, you can change one or more of your old programs to use the field level descriptions.

## Dictionary-Described Data

A program-described file can be dictionary-described. You can describe the record format information using interactive data definition utility (IDDU). Even though the file is program-described, AS/400 Query, PC Support/400, and data file utility (DFU) will use the record format description stored in the data dictionary.

An externally described file can also be dictionary-described. You can use IDDU to describe a file, then create the file using IDDU. The file created is an externally described file. You can also move into the data dictionary the file description stored in an externally described file. The system always ensures that the definitions in the data dictionary and the description stored in the externally described file are identical.

## Methods of Describing Data to the System

If you want to describe a file just to the record level, you can use the record length (RCDLEN) parameter on the Create Physical File (CRTPF) and Create Source Physical File (CRTSRCPF) commands.

If you want to describe your file to the field level, several methods can be used to describe data to the database system: IDDU, SQL/400\* commands, or data description specifications (DDS).

**Note:** Because DDS has the most options for defining data for the programmer, this guide will focus on describing database files using DDS.

### OS/400 IDDU

Physical files can be described using IDDU. You might use IDDU because it is a menu-driven, interactive method of describing data. You also might be familiar with describing data using IDDU on a System/36. In addition, IDDU allows you to describe multiple-format physical files for use with Query, PC Support/400, and DFU.

When you use IDDU to describe your files, the file definition becomes part of the Operating System/400\* (OS/400\*) data dictionary.

For more information about IDDU, see the *IDDU User's Guide*.

### Structured Query Language/400

The Structured Query Language/400 can be used to describe an AS/400 database file. SQL/400 language supports statements to describe the fields in the database file, and to create the file.

You might use the SQL/400 language because it is the IBM\* Systems Application Architecture\* database language for defining and processing database data. When database files are created using the SQL/400 language, the description of the file is automatically added to a data dictionary in the SQL collection. The data dictionary (or catalog) is then automatically maintained by the system.

For more information about the SQL/400 language, see the *SQL/400\* Programmer's Guide*.

### OS/400 DDS

Externally described data files can be described using DDS. Using DDS, you provide descriptions of the field, record, and file level information.

You might use DDS because it provides the most options for the programmer to describe data in the database. For example, only with DDS can you describe key fields in logical files.

The *DDS Coding Form* provides a common format for describing data externally. DDS data is column sensitive. The examples in this manual have numbered columns and show the data in the correct columns.

The *DDS Reference* contains a detailed description of DDS functions to describe physical and logical files.



## Naming Conventions

The file name, record format name, and field name can be as long as 10 characters and must follow all system naming conventions, but you should keep in mind that some high-level languages have more restrictive naming conventions than the system does. For example, the RPG/400\* language allows only 6-character names, while the system allows 10-character names. In some cases, you can temporarily change (rename) the system name to one that meets the high-level language restrictions. For more information about renaming database fields in programs, see your high-level language guide.

In addition, names must be unique as follows:

- Field names must be unique in a record format.
- Record format names and member names must be unique in a file.
- File names must be unique in a library.

## Describing Database Files Using DDS

When you describe a database file using DDS, you can describe information at the file, record format, join, field, key, and select/omit levels:

- File level DDS give the system information about the entire file. For example, you can specify whether all the key field values in the file must be unique.
- Record format level DDS give the system information about a specific record format in the file. For example, when you describe a logical file record format, you can specify the physical file that it is based on.
- Join level DDS give the system information about physical files used in a join logical file. For example, you can specify how to join two physical files.
- Field level DDS give the system information about individual fields in the record format. For example, you can specify the name and attributes of each field.
- Key field level DDS give the system information about the key fields for the file. For example, you can specify which fields in the record format are to be used as key fields.
- Select/omit field level DDS give the system information about which records are to be returned to the program when processing the file. Select/omit specifications apply to logical files only.

### Example of Describing a Physical File Using DDS

The DDS for a physical file must be in the following order (Figure 1-1 on page 1-6):

- 1** File level entries (optional). The UNIQUE keyword is used to indicate that the value of the key field in each record in the file must be unique. Duplicate key values are not allowed in this file.
- 2** Record format level entries. The record format name is specified, along with an optional text description.
- 3** Field level entries. The field names and field lengths are specified, along with an optional text description for each field.
- 4** Key field level entries (optional). The field names used as key fields are specified.
- 5** Comment (optional).

```

|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
A* ORDER HEADER FILE (ORDHDRP)
A 5
A
A 2 R ORDHDR 1 UNIQUE
A 3 CUST 5 0 TEXT('Customer number')
A ORDER 5 0 TEXT('Order number')
A .
A .
A .
A K CUST
A 4 K ORDER

```

Figure 1-1. DDS for a Physical File (ORDHDRP)

The following example shows a physical file ORDHDRP (an order header file), which has an arrival sequence access path without key fields specified, and the DDS necessary to describe that file.

**Record Format (ORDHDR):**

Customer Number (CUST)	Order Number (ORDER)	Order Date (ORDATE)	Purchase Order Number (CUSORD)	Shipping Instructions (SHPVIA)	Order Status (ORDSTS)	...	State (STATE)
Packed Decimal Length 5 No Decimals	Packed Decimal Length 5 No Decimals	Packed Decimal Length 6 No Decimals	Packed Decimal Length 15 No Decimals	Character Length 15	Character Length 1		Character Length 2

RSLH231-6

```

|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
A* ORDER HEADER FILE (ORDHDRP)
A R ORDHDR TEXT('Order header record')
A CUST 5 0 TEXT('Customer Number')
A ORDER 5 0 TEXT('Order Number')
A ORDATE 6 0 TEXT('Order Date')
A CUSORD 15 0 TEXT('Customer Order No.')
A SHPVIA 15 TEXT('Shipping Instr')
A ORDSTS 1 TEXT('Order Status')
A OPRNME 10 TEXT('Operator Name')
A ORDAMT 9 2 TEXT('Order Amount')
A CUTYPE 1 TEXT('Customer Type')
A INVNBR 5 0 TEXT('Invoice Number')
A PRDAT 6 0 TEXT('Printed Date')
A SEQNBR 5 0 TEXT('Sequence Number')
A OPNSTS 1 TEXT('Open Status')
A LINES 3 0 TEXT('Order Lines')
A ACTMTH 2 0 TEXT('Accounting Month')
A ACTYR 2 0 TEXT('Accounting Year')
A STATE 2 TEXT('State')
A

```

The R in position 17 indicates that a record format is being defined. The record format name ORDHDR is specified in positions 19 through 28.

You make no entry in position 17 when you are describing a field; a blank in position 17 along with a name in positions 19 through 28 indicates a field name.

The data type is specified in position 35. The valid data types are:

<b>Entry</b>	<b>Meaning</b>
A	Character
P	Packed decimal
S	Zoned decimal
B	Binary
F	Floating point
H	Hexadecimal
L	Date
T	Time
Z	Timestamp

**Notes:**

1. For double-byte character set (DBCS) data types, see Appendix B, "Double-Byte Character Set (DBCS) Considerations."
2. The AS/400 system performs arithmetic operations more efficiently for packed decimal than for zoned decimal.
3. Some high-level languages do not support floating-point data.
4. Some special considerations that apply when you are using floating-point fields are:
  - The precision associated with a floating-point field is a function of the number of bits (single or double precision) and the internal representation of the floating-point value. This translates into the number of decimal digits supported in the significant and the maximum values that can be represented in the floating-point field.
  - When a floating-point field is defined with fewer digits than supported by the precision specified, that length is only a presentation length and has no effect on the precision used for internal calculations.
  - Although floating-point numbers are accurate to 7 (single) or 15 (double) decimal digits of precision, you can specify up to 9 or 17 digits. You can use the extra digits to uniquely establish the internal bit pattern in the internal floating-point format so identical results are obtained when a floating-point number in internal format is converted to decimal and back again to internal format.

If the data type (position 35) is not specified, the decimal positions entry is used to determine the data type. If the decimal positions (positions 36 through 37) are blank, the data type is assumed to be character (A); if these positions contain a number 0 through 31, the data type is assumed to be packed decimal (P).

The length of the field is specified in positions 30 through 34, and the number of decimal positions (for numeric fields) is specified in positions 36 and 37. If a packed or zoned decimal field is to be used in a high-level language program, the field length must be limited to the length allowed by the high-level language you are using. The length is not the length of the field in storage but the number of digits or characters specified externally from storage. For example, a 5-digit packed decimal field has a length of 5 specified in DDS, but it uses only 3 bytes of storage.

Character or hexadecimal data can be defined as variable length by specifying the VARLEN field level keyword. Generally you would use variable length fields, for example, as an employee name within a database. Names usually can be stored in a 30-byte field; however, there are times when you need 100 bytes to store a very long name. If you always define the field as 100 bytes, you waste storage. If you always define the field as 30 bytes, some names are truncated.

You can use the DDS VARLEN keyword to define a character field as variable length. You can define this field as:

- Variable-length with no allocated length. This allows the field to be stored using only the number of bytes equal to the data (plus two bytes per field for the length value and a few overhead bytes per record). However, performance might be affected because all data is stored in the variable portion of the file, which requires two disk read operations to retrieve.
- Variable-length with an allocated length equal to the most likely size of the data. This allows most field data to be stored in the fixed portion of the file and minimizes unused storage allocations common with fixed-length field definitions. Only one read operation is required to retrieve field data with a length less than the allocated field length. Field data with a length greater than the allocated length is stored in the variable portion of the file and requires two read operations to retrieve the data.

### Example of Describing a Logical File Using DDS

The DDS for a logical file must be in the following order (Figure 1-2):

- 1** File level entries (optional). In this example, the UNIQUE keyword indicates that for this file the key value for each record must be unique; no duplicate key values are allowed.

For each record format:

- 2** Record format level entries. In this example, the record format name, the associated physical file, and an optional text description are specified.
- 3** Field level entries (optional). In this example, each field name used in the record format is specified.
- 4** Key field level entries (optional). In this example, the *Order* field is used as a key field.
- 5** Select/omit field level entries (optional). In this example, all records whose *Opnsts* field contains a value of N are omitted from the file's access path. That is, programs reading records from this file will never see a record whose *Opnsts* field contains an N value.
- 6** Comment.

```

|.....1.....2.....3.....4.....5.....6.....7.....8
|
A* ORDER HEADER FILE (ORDHDRP)
A   6
A
A   2 R ORDHDR
A   3 ORDER
A   CUST
A   .
A   .
A   .
A   4 K ORDER
A   O OPNSTS
A   S
A   1 UNIQUE
A   FFILE(ORDHDRP)
A   TEXT('Order number')
A   TEXT('Customer number')
A   5 CMP(EQ 'N')
A   ALL

```

Figure 1-2. DDS for a Simple Logical File (ORDHDRP)

A logical file must be created after all physical files on which it is based are created. The PFILE keyword in the previous example is used to specify the physical file or files on which the logical file is based.



Record formats in a logical file can be:

- A new record format based on fields from a physical file
- The same record format as in a previously described physical or logical file (see “Sharing Existing Record Format Descriptions” on page 1-13)

Fields in the logical file record format must either appear in the record format of at least one of the physical files or be derived from the fields of the physical files on which the logical file is based.

For more information about describing logical files, see Chapter 3, “Setting Up Logical Files.”

### **Additional Field Definition Functions**

You can describe additional information about the fields in the physical and logical file record formats with function keywords (positions 45 through 80 on the *DDS Coding Form*). Some of the things you can specify include:

- Validity checking keywords to verify that the field data meets your standards. For example, you can describe a field to have a valid range of 500 to 900. (This checking is done only when data is typed on a keyboard to the display.)
- Editing keywords to control how a field should be displayed or printed. For example, you can use the EDTCDE(Y) keyword to specify that a date field is to appear as MM/DD/YY. The EDTCDE and EDTWRD keywords can be used to control editing. (This editing is done only when used in a display or printer file.)
- Documentation, heading, and name control keywords to control the description and name of a field. For example, you can use the TEXT keyword to document a description of each field. This text description is included in your compiler list to better document the files used in your program. The TEXT and COLHDG keywords control text and column-heading definitions. The ALIAS keyword can be used to provide a more descriptive name for a field. The alias, or alternative name, is used in a program (if the high-level language supports alias names).
- Content and default value keywords to control the null content and default data for a field. The ALWNULL keyword specifies whether a null value is allowed in the field. If ALWNULL is used, the default value of the field is null. If ALWNULL is not present at the field level, the null value is not allowed, character and hexadecimal fields default to blanks, and numeric fields default to zeros, unless the DFT (default) keyword is used to specify a different value.

### **Using Existing Field Descriptions and Field Reference Files**

If a field was already described in an existing file, and you want to use that field description in a new file you are setting up, you can request the system to copy that description into your new file description. The DDS keywords REF and REFFLD allow you to refer to a field description in an existing file. This helps reduce the effort of coding DDS statements. It also helps ensure that the field attributes are used consistently in all files that use the field.

In addition, you can create a physical file for the sole purpose of using its field descriptions. That is, the file does not contain data; it is used only as a reference for the field descriptions for other files. This type of file is known as a field reference file. A **field reference file** is a physical file containing no data, just field descriptions.

You can use a field reference file to simplify record format descriptions and to ensure field descriptions are used consistently. You can define all the fields you need for an application or any group of files in a field reference file. You can create a field reference file using DDS and the Create Physical File (CRTPF) command.

After the field reference file is created, you can build physical file record formats from this file without describing the characteristics of each field in each file. When you build physical files, all you need to do is refer to the field reference file (using the REF and REFFLD keywords) and specify any changes. Any changes to the field descriptions and keywords specified in your new file override the descriptions in the field reference file.

In the following example, a field reference file named DSTREFP is created for distribution applications. Figure 1-3 on page 1-11 shows the DDS needed to describe DSTREFP.

```

| |...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
| |
| |A* FIELD REFERENCE FILE (DSTREFF)
| |A       R DSTREF                TEXT('Field reference file')
| |A
| |A* FIELDS DEFINED BY CUSTOMER MASTER RECORD (CUSMST)
| |A       CUST          5 0      TEXT('Customer numbers')
| |A                                COLHDG('CUSTOMER' 'NUMBER')
| |A       NAME          20      TEXT('Customer name')
| |A       ADDR          20      TEXT('Customer address')
| |A
| |A       CITY          20      TEXT('Customer city')
| |A
| |A       STATE         2        TEXT('State abbreviation')
| |A                                CHECK(MF)
| |A       CRECHK        1        TEXT('Credit check')
| |A                                VALUES('Y' 'N')
| |A       SEARCH        6 0      TEXT('Customer name search')
| |A                                COLHDG('SEARCH CODE')
| |A       ZIP           5 0      TEXT('Zip code')
| |A                                CHECK(MF)
| |A       CUTYPE        15      COLHDG('CUSTOMER' 'TYPE')
| |A                                RANGE(1 5)
| |A
| |A* FIELDS DEFINED BY ITEM MASTER RECORD (ITMAST)
| |A       ITEM          5        TEXT('Item number')
| |A                                COLHDG('ITEM' 'NUMBER')
| |A                                CHECK(M10)
| |A       DESCRP        18      TEXT('Item description')
| |A       PRICE         5 2      TEXT('Price per unit')
| |A                                EDTCDE(J)
| |A                                CMP(GT 0)
| |A                                COLHDG('PRICE')
| |A       ONHAND        5 0      TEXT('On hand quantity')
| |A                                EDTCDE(Z)
| |A                                CMP(GE 0)
| |A                                COLHDG('ON HAND')
| |A       WHSLOC        3        TEXT('Warehouse location')
| |A                                CHECK(MF)
| |A                                COLHDG('BIN NO')
| |A       ALLOC          R        REFFLD(ONHAND *SRC)
| |A                                TEXT('Allocated quantity')
| |A                                CMP(GE 0)
| |A                                COLHDG('ALLOCATED')
| |A
| |A* FIELDS DEFINED BY ORDER HEADER RECORD (ORDHDR)
| |A       ORDER         5 0      TEXT('Order number')
| |A                                COLHDG('ORDER' 'NUMBER')
| |A       ORDATE        6 0      TEXT('Order date')
| |A                                EDTCDE(Y)
| |A                                COLHDG('DATE' 'ORDERED')
| |A       CUSORD        15      TEXT('Cust purchase ord no.')
| |A                                COLHDG('P.O.' 'NUMBER')
| |A       SHPVIA        15      TEXT('Shipping instructions')
| |A       ORDSTS        1        TEXT('Order status code')
| |A                                COLHDG('ORDER' 'STATUS')
| |A       OPRNME        R        REFFLD(NAME *SRC)
| |A                                TEXT('Operator name')
| |A                                COLHDG('OPERATOR NAME')
| |A       ORDAMT        9 2      TEXT('Total order value')
| |A                                COLHDG('ORDER' 'AMOUNT')
| |A

```

Figure 1-3 (Part 1 of 2). DDS for a Field Reference File (DSTREFF)

	1	2	3	4	5	6	7	8
A	INVNBR	5	0		TEXT('Invoice number')			
A					COLHDG('INVOICE' 'NUMBER')			
A	PRTDAT	6	0		EDTCDE(Y)			
A					COLHDG('PRINTED' 'DATE')			
A	SEQNBR	5	0		TEXT('Sequence number')			
A					COLHDG('SEQ' 'NUMBER')			
A	OPNSTS	1			TEXT('Open status')			
A					COLHDG('OPEN' 'STATUS')			
A	LINES	3	0		TEXT('Lines on invoice')			
A					COLHDG('TOTAL' 'LINES')			
A	ACTMTH	2	0		TEXT('Accounting month')			
A					COLHDG('ACCT' 'MONTH')			
A	ACTYR	2	0		TEXT('Accounting year')			
A					COLHDG('ACCT' 'YEAR')			
A	A* FIELDS DEFINED BY ORDER DETAIL/LINE ITEM RECORD (ORDDTL)							
A	LINE	3	0		TEXT('Line no. this item')			
A					COLHDG('LINE' 'NO')			
A	QTYORD	3	0		TEXT('Quantity ordered')			
A					COLHDG('QTY' 'ORDERED')			
A					CMP(GE 0)			
A	EXTENS	6	2		TEXT('Ext of QTYORD x PRICE')			
A					EDTCDE(J)			
A					COLHDG('EXTENSION')			
A	A* FIELDS DEFINED BY ACCOUNTS RECEIVABLE							
A	ARBAL	8	2		TEXT('A/R balance due')			
A					EDTCDE(J)			
A	A* WORK AREAS AND OTHER FIELDS THAT OCCUR IN MULTIPLE PROGRAMS							
A	STATUS	12			TEXT('status description')			
A								

Figure 1-3 (Part 2 of 2). DDS for a Field Reference File (DSTREFFP)

Assume that the DDS in Figure 1-3 is entered into a source file FRSSOURCE; the member name is DSTREFFP. To then create a field reference file, use the Create Physical File (CRTPF) command as follows:

```
CRTPF FILE(DSTPRODLB/DSTREFFP)
      SRCFILE(QGPL/FRSSOURCE) MBR(*NONE)
      TEXT('Distribution field reference file')
```

The parameter MBR(\*NONE) tells the system not to add a member to the file (because the field reference file never contains data and therefore does not need a member).

To describe the physical file ORDHDRP by referring to DSTREFFP, use the following DDS (Figure 1-4 on page 1-13):

```

|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
A* ORDER HEADER FILE (ORDHDRP) - PHYSICAL FILE RECORD DEFINITION
A                                     REF(DSTREFP)
A          R ORDHDR                   TEXT('Order header record')
A          CUST                       R
A          ORDER                      R
A          ORDATE                     R
A          CUSORD                     R
A          SHPVIA                     R
A          ORDSTS                     R
A          OPRNME                     R
A          ORDAMT                     R
A          CUTYPE                     R
A          INVNBR                     R
A          PRTDAT                     R
A          SEQNBR                     R
A          OPNSTS                     R
A          LINES                      R
A          ACTMTH                     R
A          ACTYR                      R
A          STATE                      R
A

```

Figure 1-4. DDS for a Physical File (ORDHDRP) Built from a Field Reference File

The REF keyword (positions 45 through 80) with DSTREFP (the field reference file name) specified indicates the file from which field descriptions are to be used. The R in position 29 of each field indicates that the field description is to be taken from the reference file.

When you create the ORDHDRP file, the system uses the DSTREFP file to determine the attributes of the fields included in the ORDHDR record format. To create the ORDHDRP file, use the Create Physical File (CRTPF) command. Assume that the DDS in Figure 1-4 was entered into a source file QDDSSRC; the member name is ORDHDRP.

```

CRTPF FILE(DSTPRODLB/ORDHDRP)
      TEXT('Order Header physical file')

```

**Note:** The files used in some of the examples in this guide refer to this field reference file.

### Using a Data Dictionary for Field Reference

You can use a data dictionary and IDDU as an alternative to using a DDS field reference file. IDDU allows you to define fields in a data dictionary. For more information, see the *IDDU User's Guide*.

### Sharing Existing Record Format Descriptions

A record format can be described once in either a physical or a logical file (except a join logical file) and can be used by many files. When you describe a new file, you can specify that the record format of an existing file is to be used by the new file. This can help reduce the number of DDS statements that you would normally code to describe a record format in a new file and can save auxiliary storage space.

The file originally describing the record format can be deleted without affecting the files sharing the record format. After the last file using the record format is deleted, the system automatically deletes the record format description.

The following shows the DDS for two files. The first file describes a record format, and the second shares the record format of the first:

```

| .....1.....+.....2.....+.....3.....+.....4.....+.....5.....+.....6.....+.....7.....+.....8
|  A          R RECORD1                                PFILE(CUSMSTP)
|  A          CUST
|  A          NAME
|  A          ADDR
|  A          SEARCH
|  A          K CUST
|  A

```

Figure 1-5. DDS for a Logical File (CUSMSTL)

```

| .....1.....+.....2.....+.....3.....+.....4.....+.....5.....+.....6.....+.....7.....+.....8
|  A          R RECORD1                                PFILE(CUSMSTP)
|  A          K NAME                                  FORMAT(CUSMSTL)
|  A

```

Figure 1-6. DDS for a Logical File (CUSTMSTL1) Sharing a Record Format

The example shown in Figure 1-5 shows file CUSMSTL, in which the fields *Cust*, *Name*, *Addr*, and *Search* make up the record format. The *Cust* field is specified as a key field.

The DDS in Figure 1-6 shows file CUSTMSTL1, in which the **FORMAT** keyword names CUSMSTL to supply the record format. The record format name must be RECORD1, the same as the record format name shown in Figure 1-5. Because the files are sharing the same format, both files have fields *Cust*, *Name*, *Addr*, and *Search* in the record format. In file CUSMSTL1, a different key field, *Name* is specified.

The following restrictions apply to shared record formats:

- A physical file cannot share the format of a logical file.
- A join logical file cannot share the format of another file, and another file cannot share the format of a join logical file.
- A view cannot share the format of another file, and another file cannot share the format of a view. (In SQL, a **view** is an alternative representation of data from one or more tables. A view can include all or some of the columns contained in the table or tables on which it is defined.)

If the original record format is changed by deleting all related files and creating the original file and all the related files again, it is changed for all files that share it. If only the file with the original format is deleted and re-created with a new record format, all files previously sharing that file's format continue to use the original format.

If a logical file is defined but no field descriptions are specified and the **FORMAT** keyword is not specified, the record format of the first physical file (specified first on the **PFILE** keyword for the logical file) is automatically shared. The record format name specified in the logical file must be the same as the record format name specified in the physical file.

To find out if a file shares a format with another file, use the **RCDFMT** parameter on the Display Database Relations (**DSPDBR**) command.

## Describing the Access Path for the File

An access path describes the order in which records are to be retrieved. Records in a physical or logical file can be retrieved using an arrival sequence access path or a keyed sequence access path. For logical files, you can also select and omit records based on the value of one or more fields in each record.

### Arrival Sequence Access Path

The arrival sequence access path is based on the order in which the records arrive and are stored in the file. For reading or updating, records can be accessed:

- Sequentially, where each record is taken from the next sequential physical position in the file.
- Directly by relative record number, where the record is identified by its position from the start of the file.

An externally described file has an arrival sequence access path when no key fields are specified for the file.

An arrival sequence access path is valid only for the following:

- Physical files
- Logical files in which each member of the logical file is based on only one physical file member
- Join logical files
- Views

### Notes:

1. Arrival sequence is the only processing method that allows a program to use the storage space previously occupied by a deleted record by placing another record in that storage space. This method requires explicit insertion of a record given a relative record number that you provide. Another method, in which the system manages the space created by deleting records, is the reuse deleted records attribute that can be specified for physical files. For more information and tips on using the reuse deleted records attribute, see “Reusing Deleted Records” on page 5-3. For more information about processing deleted records, see “Deleting Database Records” on page 7-11.
2. Through your high-level language, the Display Physical File Member (DSPPFM) command, and the Copy File (CPYF) command, you can process a keyed sequence file in arrival sequence. You can use this function for a physical file, a simple logical file based on one physical file member, or a join logical file.
3. Through your high-level language, you can process a keyed sequence file directly by relative record number. You can use this function for a physical file, a simple logical file based on one physical file member, or a join logical file.
4. An arrival sequence access path does not take up any additional storage and is always saved or restored with the file. (Because the arrival sequence access path is nothing more than the physical order of the data as it was stored, when you save the data you save the arrival sequence access path.)

## Keyed Sequence Access Path

A keyed sequence access path is based on the contents of the key fields as defined in DDS. This type of access path is updated whenever records are added or deleted, or when records are updated and the contents of a key field is changed. The keyed sequence access path is valid for both physical and logical files. The sequence of the records in the file is defined in DDS when the file is created and is maintained automatically by the system.

Key fields defined as character fields are arranged based on the sequence defined for EBCDIC characters. Key fields defined as numeric fields are arranged based on their algebraic values, unless the DDS UNSIGNED (unsigned value) or ABSVAL (absolute value) keywords are specified for the field. Key fields defined as DBCS are allowed, but are arranged only as single bytes based on their bit representation.

**Arranging Key Fields Using an Alternative Collating Sequence:** Keyed fields that are defined as character fields can be arranged based either on the sequence for EBCDIC characters or on an alternative collating sequence. Consider the following records:

Record	Empname	Deptnbr	Empnbr
1	Jones, Mary	45	23318
2	Smith, Ron	45	41321
3	JOHNSON, JOHN	53	41322
4	Smith, ROBERT	27	56218
5	JONES, MARTIN	53	62213

If the *Empname* is the key field and is a character field, using the sequence for EBCDIC characters, the records would be arranged as follows:

Record	Empname	Deptnbr	Empnbr
1	Jones, Mary	45	23318
3	JOHNSON, JOHN	53	41322
5	JONES, MARTIN	53	62213
2	Smith, Ron	45	41321
4	Smith, ROBERT	27	56218

Notice that the EBCDIC sequence causes an unexpected sort order because the lowercase characters are sorted before uppercase characters. Thus, Smith, Ron sorts before SMITH, ROBERT. An alternative collating sequence could be used to sort the records when the records were entered using uppercase and lowercase as shown in the following example:

Record	Empname	Deptnbr	Empnbr
3	JOHNSON, JOHN	53	41322
5	JONES, MARTIN	53	62213
1	Jones, Mary	45	23318
4	Smith, ROBERT	27	56218
2	Smith, Ron	45	41321

To use an alternative collating sequence for a character key field, specify the ALTSEQ DDS keyword, and specify the name of the table containing the alternative collating sequence. When setting up a table, each 2-byte position in the table cor-



responds to a character. To change the order in which a character is sorted, change its 2-digit value to the same value as the character it should be sorted equal to. For more information about the ALTSEQ keyword, see the *DDS Reference*. For information about sorting uppercase and lowercase characters regardless of their case, the QCASE256 table in library QUSRSYS is provided for you.

**Arranging Key Fields Using the SRTSEQ Parameter:** You can arrange key fields containing character data according to several sorting sequences available with the SRTSEQ parameter. Consider the following records:

Record	Empname	Deptnbr	Empnbr
1	Jones, Marilyn	45	23318
2	Smith, Ron	45	41321
3	JOHNSON, JOHN	53	41322
4	Smith, ROBERT	27	56218
5	JONES, MARTIN	53	62213
6	Jones, Martin	08	29231

If the *Empname* field is the key field and is a character field, the \*HEX sequence (the EBCDIC sequence) arranges the records as follows:

Record	Empname	Deptnbr	Empnbr
1	Jones, Marilyn	45	23318
6	Jones, Martin	08	29231
3	JOHNSON, JOHN	53	41322
5	JONES, MARTIN	53	62213
2	Smith, Ron	45	41321
4	Smith, ROBERT	27	56218

Notice that with the \*HEX sequence, all lowercase characters are sorted before the uppercase characters. Thus, Smith, Ron sorts before SMITH, ROBERT, and JOHNSON, JOHN sorts between the lowercase and uppercase Jones. You can use the \*LANGIDSHR sort sequence to sort records when the records were entered using a mixture of uppercase and lowercase. The \*LANGIDSHR sequence, which uses the same collating weight for lowercase and uppercase characters, results in the following:

Record	Empname	Deptnbr	Empnbr
3	JOHNSON, JOHN	53	41322
1	Jones, Marilyn	45	23318
5	JONES, MARTIN	53	62213
6	Jones, Martin	08	29231
4	Smith, ROBERT	27	56218
2	Smith, Ron	45	41321

Notice that with the \*LANGIDSHR sequence, the lowercase and uppercase characters are treated as equal. Thus, JONES, MARTIN and Jones, Martin are equal and sort in the same sequence they have in the base file. While this is not incorrect, it would look better in a report if all the lowercase Jones preceded the uppercase JONES. You can use the \*LANGIDUNQ sort sequence to sort the records when the records were entered using an inconsistent uppercase and lowercase. The

\*LANGIDUNQ sequence, which uses different but sequential collating weights for lowercase and uppercase characters, results in the following:

Record	Empname	Deptnbr	Empnbr
3	JOHNSON, JOHN	53	41322
1	Jones, Marilyn	45	23318
6	Jones, Martin	08	29231
5	JONES, MARTIN	53	62213
4	Smith, ROBERT	27	56218
2	Smith, Ron	45	41321

The \*LANGIDSHR and \*LANGIDUNQ sort sequences exist for every language supported in your system. The LANGID parameter determines which \*LANGIDSHR or \*LANGIDUNQ sort sequence to use. Use the SRTSEQ parameter to specify the sort sequence and the LANGID parameter to specify the language.

**Arranging Key Fields in Ascending or Descending Sequence:** Key fields can be arranged in either ascending or descending sequence. Consider the following records:

Record	Empnbr	Clsnbr	Clsnam	Cpdate
1	56218	412	Welding I	032188
2	41322	412	Welding I	011388
3	64002	412	Welding I	011388
4	23318	412	Welding I	032188
5	41321	412	Welding I	051888
6	62213	412	Welding I	032188

If the *Empnbr* field is the key field, the two possibilities for organizing these records are:

- In ascending sequence, where the order of the records in the access path is:

Record	Empnbr	Clsnbr	Clsnam	Cpdate
4	23318	412	Welding I	032188
5	41321	412	Welding I	051888
2	41322	412	Welding I	011388
1	56218	412	Welding I	032188
6	62213	412	Welding I	032188
3	64002	412	Welding I	011388

- In descending sequence, where the order of the records in the access path is:

Record	Empnbr	Clsnbr	Clsnam	Cpdate
3	64002	412	Welding I	011388
6	62213	412	Welding I	032188
1	56218	412	Welding I	032188
2	41322	412	Welding I	011388
5	41321	412	Welding I	051888
4	23318	412	Welding I	032188

When you describe a key field, the default is ascending sequence. However, you can use the DESCEND DDS keyword to specify that you want to arrange a key field in descending sequence.

**Using More Than One Key Field:** You can use more than one key field to arrange the records in a file. The key fields do not have to use the same sequence. For example, when you use two key fields, one field can use ascending sequence while the other can use descending sequence. Consider the following records:

Record	Order	Ordate	Line	Item	Qtyord	Extens
1	52218	063088	01	88682	425	031875
2	41834	062888	03	42111	30	020550
3	41834	062888	02	61132	4	021700
4	52218	063088	02	40001	62	021700
5	41834	062888	01	00623	50	025000

If the access path uses the *Order* field, then the *Line* field as the key fields, both in ascending sequence, the order of the records in the access path is:

Record	Order	Ordate	Line	Item	Qtyord	Extens
5	41834	062888	01	00623	50	025000
3	41834	062888	02	61132	4	021700
2	41834	062888	03	42111	30	020550
1	52218	063088	01	88682	425	031875
4	52218	063088	02	40001	62	021700

If the access path uses the key field *Order* in ascending sequence, then the *Line* field in descending sequence, the order of the records in the access path is:

Record	Order	Ordate	Line	Item	Qtyord	Extens
2	41834	062888	03	42111	30	020550
3	41834	062888	02	61132	4	021700
5	41834	062888	01	00623	50	025000
4	52218	063088	02	40001	62	021700
1	52218	063088	01	88682	425	031875

When a record has key fields whose contents are the same as the key field in another record in the same file, then the file is said to have records with duplicate key values. However, the duplication must occur for *all* key fields for a record if they are to be called duplicate key values. For example, if a record format has two key fields *Order* and *Ordate*, duplicate key values occur when the contents of both the *Order* and *Ordate* fields are the same in two or more records. These records have duplicate key values:

Order	Ordate	Line	Item	Qtyord	Extens
41834	062888	03	42111	30	020550
41834	062888	02	61132	04	021700
41834	062888	01	00623	50	025000

Using the *Line* field as a third key field defines the file so that there are no duplicate keys:

(First Key Field) Order	(Second Key Field) Ordate	(Third Key Field) Line	Item	Qtyord	Extens
41834	062888	03	42111	30	020550
41834	062888	02	61132	04	021700
41834	062888	01	00623	50	025000

A logical file that has more than one record format can have records with duplicate key values, even though the record formats are based on different physical files. That is, even though the key values come from different record formats, they are considered duplicate key values.

**Preventing Duplicate Key Values:** The AS/400 database management system allows records with duplicate key values in your files. However, you may want to prevent duplicate key values in some of your files. For example, you can create a file where the key field is defined as the customer number field. In this case, you want the system to ensure that each record in the file has a unique customer number.

You can prevent duplicate key values in your files by specifying the UNIQUE keyword in DDS. With the UNIQUE keyword specified, a record cannot be entered or copied into a file if its key value is the same as the key value of a record already existing in the file.

If records with duplicate key values already exist in a physical file, the associated logical file cannot have the UNIQUE keyword specified. If you try to create a logical file with the UNIQUE keyword specified, and the associated physical file contains duplicate key values, the logical file is not created. The system sends you a message stating this and sends you messages (as many as 20) indicating which records contain duplicate key values.

When the UNIQUE keyword is specified for a file, any record added to the file cannot have a key value that duplicates the key value of an existing record in the file, regardless of the file used to add the new record. For example, two logical files LF1 and LF2 are based on the physical file PF1. The UNIQUE keyword is specified for LF1. If you use LF2 to add a record to PF1, you cannot add the record if it causes a duplicate key value in LF1.

If any of the key fields allow null values, null values that are inserted into those fields may or may not cause duplicates depending on how the access path was defined at the time the file was created. The \*INCNULL parameter of the UNIQUE keyword indicates that null values are included when determining whether duplicates exist in the unique access path. The \*EXCNULL parameter indicates that null values are not included when determining whether duplicate values exist. For more information, see the *DDS Reference*.

The following shows the DDS for a logical file that requires unique key values:

```

|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
A* ORDER TRANSACTION LOGICAL FILE (ORDFILL)
A                                     UNIQUE
A          R ORDHDR                   PFILE(ORDHDRP)
A          K ORDER
A
A          R ORDDTL                   PFILE(ORDDTLP)
A          K ORDER
A          K LINE
A

```

In this example, the contents of the key fields (the *Order* field for the ORDHDR record format, and the *Order* and *Line* fields for the ORDDTL record format) must be unique whether the record is added through the ORDHDRP file, the ORDDTLP file, or the logical file defined here. With the *Line* field specified as a second key field in the ORDDTL record format, the same value can exist in the *Order* key field in both physical files. Because the physical file ORDDTLP has two key fields and the physical file ORDHDRP has only one, the key values in the two files do not conflict.

**Arranging Duplicate Keys:** If you do not specify the UNIQUE keyword in DDS, you can specify how the system is to store records with duplicate key values, should they occur. You specify that records with duplicate key values are stored in the access path in one of the following ways:

- Last-in-first-out (LIFO). When the LIFO keyword is specified (**1**), records with duplicate key values are retrieved in last-in-first-out order by the physical sequence of the records. Below is an example of DDS using the LIFO keyword.

```

|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
A* ORDERP2
A                                     1 LIFO
A          R ORDER2
A          .
A          .
A          .
A          K ORDER
A

```

- First-in-first-out (FIFO). If the FIFO keyword is specified, records with duplicate key values are retrieved in first-in-first-out order by the physical sequence of the records.
- First-changed-first-out (FCFO). If the FCFO keyword is specified, records with duplicate key values are retrieved in first-changed-first-out order by the physical sequence of the keys.
- No specific order for duplicate key fields (the default). When the FIFO, FCFO, or LIFO keywords are not specified, no guaranteed order is specified for retrieving records with duplicate keys. No specific order for duplicate key fields allows more access path sharing, which can improve performance. For more information about access path sharing, see "Using Existing Access Paths" on page 3-13.

When a simple- or multiple-format logical file is based on more than one physical file member, records with duplicate key values are read in the order in which the files and members are specified on the DTAMBRS parameter on the Create Logical

File (CRTLF) or Add Logical File Member (ADDLFM) command. Examples of logical files with more than one record format can be found in the *DDS Reference*.

The LIFO or FIFO order of records with duplicate key values is not determined by the sequence of updates made to the contents of the key fields, but solely by the physical sequence of the records in the file member. Assume that a physical file has the FIFO keyword specified (records with duplicate keys are in first-in-first-out order), and that the following shows the order in which records were added to the file:

<b>Order Records Were Added to File</b>	<b>Key Value</b>
1	A
2	B
3	C
4	C
5	D

The sequence of the access path is (FIFO, ascending key):

<b>Record Number</b>	<b>Key Value</b>
1	A
2	B
3	C
4	C
5	D

Records 3 and 4, which have duplicate key values, are in FIFO order. That is, because record 3 was added to the file before record 4, it is read before record 4. This would become apparent if the records were read in descending order. This could be done by creating a logical file based on this physical file, with the DESCEND keyword specified in the logical file.

The sequence of the access path is (FIFO, descending key):

<b>Record Number</b>	<b>Key Value</b>
5	D
3	C
4	C
2	B
1	A

If physical record 1 is changed such that the key value is C, the sequence of the access path for the physical file is (FIFO, ascending key):

<b>Record Number</b>	<b>Key Value</b>
2	B
1	C
3	C
4	C
5	D

Finally, changing to descending order, the new sequence of the access path for the logical file is (FIFO, descending key):

Record Number	Key Value
5	D
1	C
3	C
4	C
2	B

After the change, record 1 does not appear after record 4, even though the contents of the key field were updated after record 4 was added.

The FCFO order of records with duplicate key values is determined by the sequence of updates made to the contents of the key fields. In the example above, after record 1 is changed such that the key value is C, the sequence of the access path (FCFO, ascending key only) is:

Record Number	Key Value
2	B
3	C
4	C
1	C
5	D

For FCFO, the duplicate key ordering can change when the FCFO access path is rebuilt or when a rollback operation is performed. In some cases, your key field can change but the physical key does not change. In these cases, the FCFO ordering does not change, even though the key field has changed. For example, when the index ordering is changed to be based on the absolute value of the key, the FCFO ordering does not change. The physical value of the key does not change even though your key changes from negative to positive. Because the physical key does not change, FCFO ordering does not change.

If the reuse deleted records attribute is specified for a physical file, the duplicate key ordering must be allowed to default or must be FCFO. The reuse deleted records attribute is not allowed for the physical file if either the key ordering for the file is FIFO or LIFO, or if any of the logical files defined over the physical file have duplicate key ordering of FIFO or LIFO.

### Using Existing Access Path Specifications

You can use the DDS keyword REFACCPATH to use another file's access path specifications. When the file is created, the system determines which access path to share. The file using the REFACCPATH keyword does not necessarily share the access path of the file specified in the REFACCPATH keyword. The REFACCPATH keyword is used to simply reduce the number of DDS statements that must be specified. That is, rather than code the key field specifications for the file, you can specify the REFACCPATH keyword. When the file is created, the system copies the key field and select/omit specifications from the file specified on the REFACCPATH keyword to the file being created.

## Using Floating Point Fields in Access Paths

The collating sequence for records in a keyed database file depends on the presence of the SIGNED, UNSIGNED, and ABSVAL DDS keywords. For floating-point fields, the sign is the farthest left bit, the exponent is next, and the significant is last. The collating sequence with UNSIGNED specified is:

- Positive real numbers—positive infinity
- Negative real numbers—negative infinity

A floating-point key field with the SIGNED keyword specified, or defaulted to, on the DDS has an algebraic numeric sequence. The collating sequence is negative infinity—real numbers—positive infinity.

A floating-point key field with the ABSVAL keyword specified on the DDS has an absolute value numeric sequence.

The following floating-point collating sequences are observed:

- Zero (positive or negative) collates in the same manner as any other positive/negative real number.
- Negative zero collates before positive zero for SIGNED sequences.
- Negative and positive zero collate the same for ABSVAL sequences.

You cannot use not-a-number (\*NAN) values in key fields. If you attempt this, and a \*NAN value is detected in a key field during file creation, the file is not created.

---

## Database File Creation: Introduction

The system supports several methods for creating a database file:

- OS/400 IDDU
- Structured Query Language/400
- OS/400 control language (CL)

You can create a database file using IDDU. If you are using IDDU to describe your database files, you might also consider using it to create your files.

You can create a database file using the SQL/400 statements. SQL/400 language is the IBM Systems Application Architecture\* (SAA\*) relational database language, and can be used on the AS/400 system to interactively describe and create database files.

You can also create a database file using CL. The CL database file create commands are: Create Physical File (CRTPF), Create Logical File (CRTLF), and Create Source Physical File (CRTSRCPF).

Because additional system functions are available with CL, this guide focuses on creating files using CL.



---

## Database File and Member Attributes: Introduction

When you create a database file, database attributes are stored with the file and members. You specify attributes with database command parameters. For discussions on specifying these attributes and their possible values, see the CRTPF, CRTLF, CRTSRCPF, ADDPFM, ADDLFM, CHGPF, CHGLF, CHGPFM, CHGSRCPF, and CHGLFM commands in the *CL Reference* manual.

## File Name and Member Name (FILE and MBR) Parameters

You name a file with the FILE parameter in the create command. You also name the library in which the file will reside. When you create a physical or logical file, the system normally creates a member with the same name as the file. You can, however, specify a member name with the MBR parameter in the create commands. You can also choose not to create any members by specifying MBR(\*NONE) in the create command.

**Note:** The system does *not* automatically create a member for a source physical file.

---

## Physical File Member Control (DTAMBRS) Parameter

You can control the reading of the physical file members with the DTAMBRS parameter of the Create Logical File (CRTLF) command. You can specify:

- The order in which the physical file members are to be read.
- The number of physical file members to be used.

For more information about using logical files in this way, see “Logical File Members” on page 3-20.

## Source File and Source Member (SRCFILE and SRCMBR) Parameters

The SRCFILE and SRCMBR parameters specify the names of the source file and members containing the DDS statements that describe the file being created. If you do not specify a name:

- The default source file name is QDDSSRC.
- The default member name is the name specified on the FILE parameter.

## Database File Type (FILETYPE) Parameter

A database file type is either data (\*DATA) or source (\*SRC). The Create Physical File (CRTPF) and Create Logical File (CRTLF) commands use the default data file type (\*DATA).

## Maximum Number of Members Allowed (MAXMBRS) Parameter

The MAXMBRS parameter specifies the maximum number of members the file can hold. The default maximum number of members for physical and logical files is one, and the default for source physical files is \*NOMAX.

## Where to Store the Data (UNIT) Parameter

The system finds a place for the file on auxiliary storage. To specify where to store the file, use the UNIT parameter. The UNIT parameter specifies:

- The location of data records in physical files.
- The access path for both physical files and logical files.

The data is placed on different units if:

- There is not enough space on the unit.
- The unit is not valid for your system.

An informational message indicating that the file was not placed on the requested unit is sent when file members are added. (A message is *not* sent when the file member is extended.)

### Unit Parameter Tips

In general, you should not specify the UNIT parameter. Let the system place the file on the disk unit of its choosing. This is usually better for performance, and relieves you of the task of managing auxiliary storage.

If you specify a unit number and also an auxiliary storage pool, the unit number is ignored. For more information about auxiliary storage pools, see the *Advanced Backup and Recovery Guide*.

## Frequency of Writing Data to Auxiliary Storage (FRCRATIO) Parameter

You can control when database changes are written to auxiliary storage using the force write ratio (FRCRATIO) parameter on either the create, change, or override database file commands. Normally, the system determines when to write changed data from main storage to auxiliary storage. Closing the file (except for a shared close) and the force-end-of-data operation forces remaining updates, deletions, and additions to auxiliary storage. If you are journaling the file, the FRCRATIO parameter should normally be \*NONE.

### FRCRATIO Parameter Tip

Using the FRCRATIO parameter has performance and recovery considerations for your system. To understand these considerations, see Chapter 13, "Database Recovery Considerations."

## Frequency of Writing the Access Path (FRCACCPH) Parameter

The force access path (FRCACCPH) parameter controls when an access path is written to auxiliary storage. FRCACCPH(\*YES) forces the access path to auxiliary storage whenever the access path is changed. This reduces the chance that the access path will need to be rebuilt should the system fail.

### FRCACCPH Parameter Tips

Specifying FRCACCPH(\*YES) can degrade performance when changes occur to the access path. An alternative to forcing the access path is journaling the access path. For more information about forcing access paths and journaling access paths, see Chapter 13, "Database Recovery Considerations."

## Check for Record Format Description Changes (LVLCHK) Parameter

When the file is opened, the system checks for changes to the database file definition. When the file changes to an extent that your program may not be able to process the file, the system notifies your program. The default is to do level checking. You can specify if you want level checking when you:

- Create a file.
- Use a change database file command.

You can override the system and ignore the level check using the Override with Database File (OVRDBF) command.

### Level Check Example

For example, assume you compiled your program two months ago and, at that time, the file the program was defined as having three fields in each record. Last week another programmer decided to add a new field to the record format, so that now each record would have four fields. The system notifies your program, when it tries to open the file, that a significant change occurred to the definition of the file since the last time the program was compiled. This notification is known as a record format level check.

## Current Access Path Maintenance (MAINT) Parameter

The MAINT parameter specifies how access paths are maintained for closed files. While a file is open, the system maintains the access paths as changes are made to the data in the file. However, because more than one access path can exist for the same data, changing data in one file might cause changes to be made in access paths for other files that are not currently open (in use). The three ways of maintaining access paths of closed files are:

- **Immediate** maintenance of an access path means that the access path is maintained as changes are made to its associated data, regardless if the file is open.
- **Rebuild** maintenance of an access path means that the access path is only maintained while the file is open, not when the file is closed; the access path is rebuilt when the file is opened the next time. When a file with rebuild maintenance is closed, the system stops maintaining the access path. When the file is opened again, the access path is totally rebuilt. If one or more programs has opened a specific file member with rebuild maintenance specified, the system maintains the access path for that member until the last user closes the file member.
- **Delayed** maintenance of an access path means that any maintenance for the access path is done after the file member is opened the next time and while it remains open. However, the access path is not rebuilt as it is with rebuild maintenance. Updates to the access path are collected from the time the member is closed until it is opened again. When it is opened, only the collected changes are merged into the access path.

If you do not specify the type of maintenance for a file, the default is immediate maintenance.

## MAINT Parameter Comparison

Figure 1-7 compares immediate, rebuild, and delayed maintenance as they affect opening and processing files.

Figure 1-7. MAINT Values

Function	Immediate Maintenance	Rebuild Maintenance	Delayed Maintenance
Open	Fast open because the access path is current.	Slow open because access path must be rebuilt.	Moderately fast open because the access path does not have to be rebuilt, but it must still be changed. Slow open if extensive changes are needed.
Process	Slower update/output operations when many access paths with immediate maintenance are built over changing data (the system must maintain the access paths).	Faster update/output operations when many access paths with rebuild maintenance are built over changing data and are not open (the system does not have to maintain the access paths).	Moderately fast update/output operations when many access paths with delayed maintenance are built over changing data and are not open, (the system records the changes, but the access path itself is not maintained).

**Note:**

1. Delayed or rebuild maintenance cannot be specified for a file that has unique keys.
2. Rebuild maintenance cannot be specified for a file if its access path is being journaled.

## MAINT Parameter Tips

The type of access path maintenance to specify depends on the number of records and the frequency of additions, deletions, and updates to a file while the file is closed.

You should use delayed maintenance for files that have relatively few changes to the access path while the file members are closed. Delayed maintenance reduces system overhead by reducing the number of access paths that are maintained immediately. It may also result in faster open processing, because the access paths do not have to be rebuilt.

You may want to specify immediate maintenance for access paths that are used frequently, or when you cannot wait for an access path to be rebuilt when the file is opened. You may want to specify delayed maintenance for access paths that are not used frequently, if infrequent changes are made to the record keys that make up the access path.

In general, for files used interactively, immediate maintenance results in good response time. For files used in batch jobs, either immediate, delayed, or rebuild maintenance is adequate, depending on the size of the members and the frequency of changes.

## Access Path Recovery (RECOVER) Parameter

After a failure, changed access paths that were not forced to auxiliary storage or journaled cannot be used until they are rebuilt. The RECOVER parameter on the Create Physical File (CRTPF), the Create Logical File (CRTLF), and the Create Source Physical File (CRTSRCPF) commands specifies when that access path is to be rebuilt. Access paths are rebuilt either during the initial program load (IPL), after the IPL, or when a file is opened.

Figure 1-8 shows your choices for possible combinations of duplicate key and maintenance options.

Figure 1-8. Recovery Options

With This Duplicate Key Option	And This Maintenance Option	Your Recovery Options Are
Unique	Immediate	Rebuild during the IPL (*IPL) Rebuild after the IPL (*AFTIPL, default) Do not rebuild at IPL, wait for first open (*NO)
Not unique	Immediate or delayed	Rebuild during the IPL (*IPL) Rebuild after the IPL (*AFTIPL) Do not rebuild at IPL, wait for first open (*NO, default)
Not unique	Rebuild	Do not rebuild at IPL, wait for first open (*NO, default)

### RECOVER Parameter Tip

A list of files that have access paths that need to be recovered is shown on the Edit Rebuild of Access Paths display during the next initial program load (IPL) if the IPL is in manual mode. You can edit the original recovery option for the file by selecting the desired option on the display. After the IPL is complete, you can use the Edit Rebuild of Access Paths (EDTRBDAP) command to set the sequence in which access paths are rebuilt. If the IPL is unattended, the Edit Rebuild of Access Paths display is not shown and the access paths are rebuilt in the order determined by the RECOVER parameter. You only see the \*AFTIPL and \*NO (open) access paths.

## File Sharing (SHARE) Parameter

The database system lets multiple users access and change a file at the same time. The SHARE parameter allows sharing of opened files in the same job. For example, sharing a file in a job allows programs in the job to share a file's status, record position, and buffer. Sharing files in a job can improve performance by reducing:

- The amount of storage the job needs.
- The time required to open and close the file.

For more information about sharing files in the same job, see "Sharing Database Files in the Same Job or Activation Group" on page 5-8.

## | **Locked File or Record Wait Time (WAITFILE and WAITRCD) Parameters**

When you create a file, you can specify how long a program should wait for either the file or a record in the file if another job has the file or record locked. If the wait time ends before the file or record is released, a message is sent to the program indicating that the job was not able to use the file or read the record. For more information about record and file locks and wait times, see “Record Locks” on page 5-7 and “File Locks” on page 5-8.

## | **Public Authority (AUT) Parameter**

When you create a file, you can specify public authority. Public authority is the authority a user has to a file (or other object on the system) if that user does not have specific authority for the file or does not belong to a group with specific authority for the file. For more information about public authority, see “Public Authority” on page 4-3.

## | **System on Which the File Is Created (SYSTEM) Parameter**

You can specify if the file is to be created on the local system or a remote system that supports distributed data management (DDM). For more information about DDM, see the *DDM Guide*.

## | **File and Member Text (TEXT) Parameter**

You can specify a text description for each file and member you create. The text data is useful in describing information about your file and members.

## | **Coded Character Set Identifier (CCSID) Parameter**

You can specify a coded character set identifier (CCSID) for physical files. The CCSID describes the encoding scheme and the character set for character type fields contained in this file. For more information about CCSIDs, see the *National Language Support Planning Guide*.

## | **Sort Sequence (SRTSEQ) Parameter**

You can specify the sort sequence for a file. The values of the SRTSEQ parameter along with the CCSID and LANGID parameters determine which sort sequence table the file uses. You can set the SETSEQ parameter for both the physical and the logical files.

You can specify:

- System supplied sort sequence tables with unique or shared collating weights. There are sort sequence tables for each supported language.
- Any user-created sort sequence table.
- The hexadecimal value of the characters in the character set.
- The sort sequence of the current job or the one specified in the ALTSEQ parameter.

The sort sequence table is stored with the file, except when the sort sequence is \*HEX.

## **Language Identifier (LANGID) Parameter**

You can specify the language identifier that the system should use when the SRTSEQ parameter value is \*LANGIDSHR or \*LANGIDUNQ. The values of the LANGID, CCSID, and SRTSEQ parameters determine which sort sequence table the file uses. You can set the LANGID parameter for physical and logical files.

You can specify any language identifier supported on your system, or you can specify that the language identifier for the current job be used.





---

## Chapter 2. Setting Up Physical Files

This chapter discusses some of the unique considerations for describing, then creating, a physical file.

For information about describing a physical file record format, see “Example of Describing a Physical File Using DDS” on page 1-5.

For information about describing a physical file access path, refer to “Describing the Access Path for the File” on page 1-15.

---

### Creating a Physical File

To create a physical file, take the following steps:

1. If you are using DDS, enter DDS for the physical file into a source file. This can be done using the AS/400 Application Development Tools source entry utility (SEU). See “Working with Source Files” on page 14-3, for more information about how source statements are entered in source files.
2. Create the physical file. You can use the Create Physical File (CRTPF) command, or the Create Source Physical File (CRTSRCPF) command.

The following command creates a one-member file using DDS and places it in a library called DSTPRODLB.

```
CRTPF FILE(DSTPRODLB/ORDHDRP)
      TEXT('Order header physical file')
```

As shown, this command uses defaults. For the SRCFILE and SRCMBR parameters, the system uses DDS in the source file called QDDSSRC and the member named ORDHDRP (the same as the file name). The file ORDHDRP with one member of the same name is placed in the library DSTPRODLB.

---

### Specifying Physical File and Member Attributes

Some of the attributes you can specify for physical files and members on the Create Physical File (CRTPF), Create Source Physical File (CRTSRCPF), Change Physical File (CHGPF), Change Source Physical File (CHGSRCPF), Add Physical File Member (ADDPFM), and Change Physical File Member (CHGPFM) commands include (command names are given in parentheses):

#### Expiration Date

**EXPDATE Parameter.** This parameter specifies an expiration date for each member in the file (ADDPFM, CHGPFM, CRTPF, CHGPF, CRTSRCPF, and CHGSRCPF commands). If the expiration date is past, the system operator is notified when the file is opened. The system operator can then override the expiration date and continue, or stop the job. Each member can have a different expiration date, which is specified when the member is added to the file. (The expiration date check can be overridden; see “Checking for the File’s Expiration Date” on page 5-6.)

## Size of the Physical File Member

**SIZE Parameter.** This parameter specifies the maximum number of records that can be placed in each member (CRTPF, CHGPF, CRTSRCPF, AND CHGSRCPF commands). The following formula can be used to determine the maximum:

$$R + (I * N)$$

where:

R is the starting record count  
I is the number of records (increment) to add each time  
N is the number of times to add the increment

The defaults for the SIZE parameter are:

R 10,000  
I 1,000  
N 3 (CRTPF command)  
499 (CRTSRCPF command)

For example, assume that R is a file created for 5000 records plus 3 increments of 1000 records each. The system can add 1000 to the initial record count of 5000 three times to make the total maximum 8000. When the total maximum is reached, the system operator either stops the job or tells the system to add another increment of records and continue. When increments are added, a message is sent to the system history log. When the file is extended beyond its maximum size, the minimum extension is 10% of the current size, even if this is larger than the specified increment.

Instead of taking the default size or specifying a size, you can specify \*NOMAX. For information about the maximum number of records allowed in a file, see Appendix A, "Database File Sizes."

## Storage Allocation

**ALLOCATE Parameter.** This parameter controls the storage allocated for members when they are added to the file (CRTPF, CHGPF, CRTSRCPF, and CHGSRCPF commands). The storage allocated would be large enough to contain the initial record count for a member. If you do not allocate storage when the members are added, the system will automatically extend the storage allocation as needed. You can use the ALLOCATE parameter only if you specified a maximum size on the SIZE parameter. If SIZE(\*NOMAX) is specified, then ALLOCATE(\*YES) cannot be specified.

## Method of Allocating Storage

**CONTIG Parameter.** This parameter controls the method of allocating physical storage for a member (CRTPF and CRTSRCPF commands). If you allocate storage, you can request that the storage for the starting record count for a member be contiguous. That is, all the records in a member are to physically reside together. If there is not enough contiguous storage, contiguous storage allocation is not used and an informational message is sent to the job that requests the allocation, at the time the member is added.

**Note:** When a physical file is first created, the system always tries to allocate its initial storage contiguously. The only difference between using

CONTIG(\*NO) and CONTIG(\*YES) is that with CONTIG(\*YES) the system sends a message to the job log if it is unable to allocate contiguous storage when the file is created. No message is sent when a file is extended after creation, regardless of what you specified on the CONTIG parameter.

## Record Length

**RCDLEN Parameter.** This parameter specifies the length of records in the file (CRTPF and CRTSRCPF commands). If the file is described to the record level only, then you specify the RCDLEN parameter when the file is created. This parameter cannot be specified if the file is described using DDS, IDDU, or SQL/400 (the system automatically determines the length of records in the file from the field level descriptions).

## Deleted Records

**DLTPCT Parameter.** This parameter specifies the percentage of deleted records a file can contain before you want the system to send a message to the system history log (CRTPF, CHGPF, CRTSRCPF, and CHGSRCPF commands). When a file is closed, the system checks the member to determine the percentage of deleted records. If the percentage exceeds that value specified in the DLTPCT parameter, a message is sent to the history log. (For information about processing the history log, see the chapter on message handling in the *CL Programmer's Guide*.) One reason you might want to know when a file reaches a certain percentage of deleted records is to reclaim the space used by the deleted records. After you receive the message about deleted records, you could run the Reorganize Physical File Member (RGZPFM) command to reclaim the space. (For more information about RGZPFM, see "Reorganizing Data in Physical File Members" on page 10-3.) You can also specify to bypass the deleted records check by using the \*NONE value for the DLTPCT parameter. \*NONE is the default for the DLTPCT parameter.

**REUSEDLT Parameter.** This parameter specifies whether deleted record space should be reused on subsequent write operations (CRTPF and CHGPF commands). When you specify \*YES for the REUSEDLT parameter, all insert requests on that file try to reuse deleted record space. Reusing deleted record space allows you to reclaim space used by deleted records without having to issue a RGZPFM command. When the CHGPF command is used to change a file to reuse deleted records, the command could take a long time to run, especially if the file is large and there are already a lot of deleted records in it. It is important to note the following:

- The term *arrival order* loses its meaning for a file that reuses deleted record space. Records are no longer always inserted at the end of the file when deleted record space is reused.
- If a new physical file is created with the reuse deleted record space attribute and the file is keyed, the FIFO or LIFO access path attribute cannot be specified for the physical file, nor can any keyed logical file with the FIFO or LIFO access path attribute be built over the physical file.
- You cannot change an existing physical file to reuse deleted record space if there are any logical files over the physical file that specify FIFO or LIFO ordering for duplicate keys, or if the physical file has a FIFO or LIFO duplicate key ordering.

- Reusing deleted record space should not be specified for a file that is processed as a direct file or if the file is processed using relative record numbers.

**Note:** See “Reusing Deleted Records” on page 5-3 for more information on reusing deleted records.

\*NO is the default for the REUSEDLT parameter.

## Physical File Capabilities

**ALWUPD and ALWDLT Parameters.** File capabilities are used to control which input/output operations are allowed for a database file independent of database file authority. For more information about database file capabilities and authority, see Chapter 4, “Database Security.”

## Source Type

**SRCTYPE Parameter.** This parameter specifies the source type for a member in a source file (ADDPFM and CHGPFM commands). The source type determines the syntax checker, prompting, and formatting that are used for the member. If the user specifies a unique source type (other than AS/400 supported types like COBOL and RPG), the user must provide the programming to handle the unique type.

If the source type is changed, it is only reflected when the member is subsequently opened; members currently open are not affected.

## Chapter 3. Setting Up Logical Files

This chapter discusses some of the unique considerations for describing, then creating, a logical file. Many of the rules for setting up logical files apply to all categories of logical files. In this guide, rules that apply only to one category of logical file identify which category they refer to. Rules that apply to all categories of logical files do not identify the specific categories they apply to.

### Describing Logical File Record Formats

For every logical file record format described with DDS, you must specify a record format name and either the PFILE keyword (for simple and multiple format logical files), or the JFILE keyword (for join logical files). The file names specified on the PFILE or JFILE keyword are the physical files that the logical file is based on. A simple or multiple-format logical file record format can be specified with DDS in any one of the following ways:

1. In the simple logical file record format, specify only the record format name and the PFILE keyword. The record format for the only (or first) physical file specified on the PFILE keyword is the record format for the logical file. The record format name specified in the logical file must be the same as the record format name in the only (or first) physical file.

```
| .....1.....2.....3.....4.....5.....6.....7.....8  
| A  
| A          R ORDDTL          PFILE(ORDDTLP)  
| A
```

Figure 3-1. Simple Logical File

2. In the following example, you describe your own record format by listing the field names you want to include. You can specify the field names in a different order, rename fields using the RENAME keyword, combine fields using the CONCAT keyword, and use specific positions of a field using the SST keyword. You can also override attributes of the fields by specifying different attributes in the logical file.

```
| .....1.....2.....3.....4.....5.....6.....7.....8  
| A  
| A          R ORDHDR          PFILE(ORDHDRP)  
| A          ORDER  
| A          CUST  
| A          SHPVIA  
| A
```

Figure 3-2. Simple Logical File with Fields Specified

3. In the following example, the file name specified on the FORMAT keyword is the name of a database file. The record format is shared from this database file by the logical file being described. The file name can be qualified by a library name. If a library name is not specified, the library list is used to find the file. The file must exist when the file you are describing is created. In addition, the record format name you specify in the logical file must be the same as one of the record format names in the file you specify on the FORMAT keyword.

```

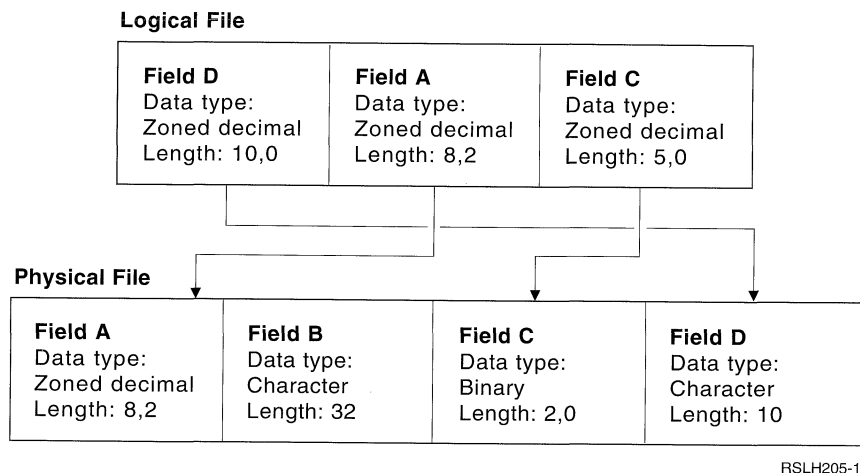
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A
  A          R CUSRCD                      PFILE(CUSMSTP)
  A                      FORMAT(CUSMSTL)
  A

```

In the following example, a program needs:

- The fields placed in a different order
- A subset of the fields from the physical file
- The data types changed for some fields
- The field lengths changed for some fields

You can use a logical file to make these changes.



For the logical file, the DDS would be:

```

|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A
  A          R LOGREC                      PFILE(PF1)
  A          D          10S 0
  A          A
  A          C          5S 0
  A

```

For the physical file, the DDS would be:

```

|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A
  A          R PHYREC
  A          A          8S 2
  A          B          32
  A          C          2B 0
  A          D          10
  A

```

When a record is read from the logical file, the fields from the physical file are changed to match the logical file description. If the program updates or adds a record, the fields are changed back. For an add or update operation using a logical file, the program must supply data that conforms with the format used by the logical file.

The following chart shows what types of data mapping are valid between physical and logical files.

Physical File Data Type	Logical File Data Type							
	Character or Hexadecimal	Zoned	Packed	Binary	Floating Point	Date	Time	Timestamp
Character or Hexadecimal	Valid	See Note 1	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid
Zoned	See Note 1	Valid	Valid	See Note 2	Valid	Not valid	Not valid	Not Valid
Packed	Not valid	Valid	Valid	See Note 2	Valid	Not valid	Not valid	Not valid
Binary	Not valid	See Note 2	See Note 2	See Note 3	See Note 2	Not valid	Not valid	Not valid
Floating Point	Not valid	Valid	Valid	See Note 2	Valid	Not valid	Not valid	Not valid
Date	Not valid	Valid	Not valid	Not valid	Not valid	Valid	Not valid	Not valid
Time	Not valid	Valid	Not valid	Not valid	Not valid	Not valid	Valid	Not valid
Time Stamp	Not valid	Not valid	Not valid	Not valid	Not valid	Valid	Valid	Valid

**Notes:**

1. Valid only if the number of characters or bytes equals the number of digits.
2. Valid only if the binary field has zero decimal positions.
3. Valid only if both binary fields have the same number of decimal positions.

**Note:** For information about mapping DBCS fields, see Appendix B, “Double-Byte Character Set (DBCS) Considerations.”

## Describing Field Use for Logical Files

You can specify that fields in database files are to be input-only, both (input/output), or neither fields. Do this by specifying one of the following in position 38:

Entry	Meaning
Blank	For simple or multiple format logical files, defaults to B (both) For join logical files, defaults to I (input only)
B	Both input and output allowed; not valid for join logical files
I	Input only (read only)
N	Neither input nor output; valid only for join logical files

**Note:** The usage value (in position 38) is not used on a reference function. When another file refers to a field (using a REF or REFFLD keyword) in a logical file, the usage value is not copied into that file.

### Both

A both field can be used for both input and output operations. Your program can read data from the field and write data to the field. Both fields are not valid for join logical files, because join logical files are read-only files.

### Input Only

An input only field can be used for read operations only. Your program can read data from the field, but cannot update the field in the file. Typical cases of input-only fields are key fields (to reduce maintenance of access paths by preventing changes to key field values), sensitive fields that a user can see but not update (for example, salary), and fields for which either the translation table (TRNTBL) keyword or the substring (SST) keyword is specified.

If your program updates a record in which you have specified input-only fields, the input-only fields are not changed in the file. If your program adds a record that has input-only fields, the input-only fields take default values (DFT keyword).

### **Neither**

A neither field is used neither for input nor for output. It is valid only for join logical files. A neither field can be used as a join field in a join logical file, but your program cannot read or update a neither field.

Use neither fields when the attributes of join fields in the physical files do not match. In this case, one or both join fields must be defined again. However, you cannot include these redefined fields in the record format (the application program does not see the redefined fields.) Therefore, redefined join fields can be coded N so that they do not appear in the record format.

A field with N in position 38 does not appear in the buffer used by your program. However, the field description is displayed with the Display File Field Description (DSPFFD) command.

Neither fields cannot be used as select/omit or key fields.

For an example of a neither field, see “Describing Fields That Never Appear in the Record Format (Example 5)” on page 3-37.

## **Deriving New Fields from Existing Fields**

Fields in a logical file can be derived from fields in the physical file the logical file is based on or from fields in the same logical file. For example, you can concatenate, using the CONCAT keyword, two or more fields from a physical file to make them appear as one field in the logical file. Likewise, you can divide one field in the physical file to make it appear as multiple fields in the logical file with the SST keyword.

### **Concatenated Fields**

Using the CONCAT keyword, you can combine two or more fields from a physical file record format to make one field in a logical file record format. For example, a physical file record format contains the fields *Month*, *Day*, and *Year*. For a logical file, you concatenate these fields into one field, *Date*.

The field length for the resulting concatenated field is the sum of the lengths of the included fields (unless the fields in the physical file are binary or packed decimal, in which case they are changed to zoned decimal). The field length of the resulting field is automatically calculated by the system. A concatenated field can have:

- Column headings
- Validity checking
- Text description
- Edit code or edit word (numeric concatenated fields only)

**Note:** This editing and validity checking information is not used by the database management system but is retrieved when field descriptions from the database file are referred to in a display or printer file.

When fields are concatenated, the data types can change (the resulting data type is automatically determined by the system). The following rules and restrictions apply:





updated record would contain 103188. If the *Date* field were specified first, the updated record would contain 123188.

Concatenated fields can also be used as key fields and select/omit fields.

### Substring Fields

You can use the SST keyword to specify which fields (character, hexadecimal, or zoned decimal) are in a substring. (You can also use substring with a packed field in a physical file by specifying S (zoned decimal) as the data type in the logical file.) For example, assume you defined the *Date* field in physical file *PF1* as 6 characters in length. You can describe the logical file with three fields, each 2 characters in length. You can use the SST keyword to define MM as 2 characters starting in position 1 of the *Date* field, DD as 2 characters starting in position 3 of the *Date* field, and YY as 2 characters starting in position 5 of the *Date* field.

The following shows the field descriptions in DDS for these substring fields. The SST keyword is used to specify the field to substring.

```

|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
|
| A          R REC1                                PFILE(PF1)
|
| A
|
| A          MM          I          SST( DATE 1 2)
|
| A          DD          I          SST( DATE 3 2)
|
| A          YY          I          SST( DATE 5 2)
|
| A

```

Note that the starting position of the substring is specified according to its position in the field being operated on (*Date*), not according to its position in the file. The I in the Usage column indicates input-only.

Substring fields can also be used as key fields and select/omit fields.

### Renamed Fields

You can name a field in a logical file differently than in a physical file using the RENAME keyword. You might want to rename a field in a logical file because the program was written using a different field name or because the original field name does not conform to the naming restrictions of the high-level language you are using.

### Translated Fields

You can specify a translation table for a field using the TRNTBL keyword. When you read a logical file record and a translation table was specified for one or more fields in the logical file, the system translates the data from the field value in the physical file to the value determined by the translation table.

## Describing Floating-Point Fields in Logical Files

You can use floating-point fields as mapped fields in logical files. A single- or double-precision floating-point field can be mapped to or from a zoned, packed, zero-precision binary, or another floating-point field. You cannot map between a floating-point field and a nonzero-precision binary field, a character field, a hexadecimal field, or a DBCS field.

Mapping between floating-point fields of different precision, single or double, or between floating-point fields and other numeric fields, can result in rounding or a

loss of precision. Mapping a double-precision floating-point number to a single-precision floating-point number can result in rounding, depending on the particular number involved and its internal representation. Rounding is to the nearest (even) bit. The result always contains as much precision as possible. A loss of precision can also occur between two decimal numbers if the number of digits of precision is decreased.

You can inadvertently change the value of a field which your program did not explicitly change. For floating-point fields, this can occur if a physical file has a double-precision field that is mapped to a single-precision field in a logical file, and you issue an update for the record through the logical file. If the internal representation of the floating-point number causes it to be rounded when it is mapped to the logical file, then the update of the logical record causes a permanent loss of precision in the physical file. If the rounded number is the key of the physical record, then the sequence of records in the physical file can also change.

A fixed-point numeric field can also be updated inadvertently if the precision is decreased in the logical file.

---

## Describing Access Paths for Logical Files

The access path for a logical file record format can be specified in one of three ways:

1. Keyed sequence access path specification. Specify key fields after the last record or field level specification. The key field names must be in the record format. For join logical files, the key fields must come from the first, or primary, physical file.

```
|
| |...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
| |   A           R CUSRCD                               PFILE(CUSMSTP)
| |   A           K ARBAL
| |   A           K CRDLMT
| |   A
```

2. Arrival sequence access path specification. Specify no key fields. You can specify only one physical file on the PFILE keyword (and only one of the physical file's members when you add the logical file member).

```
|
| |...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
| |   A           R CUSRCD                               PFILE(CUSMSTP)
```

3. Previously defined keyed-sequence access path specification (for simple and multiple format logical files only). Specify the REFACPTH keyword at the file level to identify a previously created database file whose access path and select/omit specifications are to be copied to this logical file. You cannot specify individual key or select/omit fields with the REFACPTH keyword.

**Note:** Even though the specified file's access path specifications are used, the system determines which file's access path, if any, will actually be shared. The system always tries to share access paths, regardless of whether the REFACPTH keyword is used.

```

|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
      A          R CUSRCD          REFACCPH(DSTPRODLIB/ORDHDRL)
                                   PFILE(CUSMSTP)

```

When you define a record format for a logical file that shares key field specifications of another file's access path (using the DDS keyword, REFACCPH), you can use any fields from the associated physical file record format. These fields do not have to be used in the file that describes the access path. However, all key and select/omit fields used in the file that describes the access path must be used in the new record format.

## Selecting and Omitting Records Using Logical Files

The system can select and omit records when using a logical file. This can help you to exclude records in a file for processing convenience or for security.

The process of selecting and omitting records is based on comparisons identified in position 17 of the *DDS Coding Form* for the logical file, and is similar to a series of comparisons coded in a high-level language program. For example, in a logical file that contains order detail records, you can specify that the only records you want to use are those in which the quantity ordered is greater than the quantity shipped. All other records are omitted from the access path. The omitted records remain in the physical file but are not retrieved for the logical file. If you are adding records to the physical file, all records are added, but only selected records that match the select/omit criteria can be retrieved using the select/omit access path.

In DDS, to specify select or omit, you specify an S (select) or O (omit) in position 17 of the *DDS Coding Form*. You then name the field (in positions 19 through 28) that will be used in the selection or omission process. In positions 45 through 80 you specify the comparison.

**Note:** Select/omit specifications appear after key specifications (if keys are specified).

Records can be selected and omitted by several types of comparisons:

- **VALUES.** The contents of the field are compared to a list of not more than 100 values. If a match is found, the record is selected or omitted. In the following example, a record is selected if one of the values specified in the VALUES keyword is found in the *Itmnbr* field.

```

|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
      A          S ITMNR          VALUES(301542 306902 382101 422109 +
      A          431652 486592 502356 556608 590307)
      A

```

- **RANGE.** The contents of the field are compared to lower and upper limits. If the contents are greater than or equal to the lower limit and less than or equal to the upper limit, the record is selected or omitted. In the following example, all records with a range 301000 through 599999 in the *Itmnbr* field are selected.

```

|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
      A          S ITMNR          RANGE(301000 599999)
      A

```

- **CMP.** The contents of a field are compared to a value or the contents of another field. Valid comparison codes are EQ, NE, LT, NL, GT, NG, LE, and GE. If the comparison is met, the record is selected or omitted. In the following example, a record is selected if its *Itmnbr* field is less than or equal to 599999:

```
|
| |...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
| |   A           S ITMNR                      CMP(LE 599999)
| |   A
|
```

The value for a numeric field for which the **CMP**, **VALUES**, or **RANGE** keyword is specified is aligned based on the decimal positions specified for the field and filled with zeros where necessary. If decimal positions were not specified for the field, the decimal point is placed to the right of the farthest right digit in the value. For example, for a numeric field with length 5 and decimal position 2, the value 1.2 is interpreted as 001.20 and the value 100 is interpreted as 100.00.

The status of a record is determined by evaluating select/omit statements in the sequence you specify them. If a record qualifies for selection or omission, subsequent statements are ignored.

Normally the select and omit comparisons are treated independently from one another; the comparisons are **ORed** together. That is, if the select or omit comparison is met, the record is either selected or omitted. If the condition is not met, the system proceeds to the next comparison. To connect comparisons together, you simply leave a space in position 17 of the *DDS Coding Form*. Then, all the comparisons that were connected in this fashion must be met before the record is selected or omitted. That is, the comparisons are **ANDed** together.

The fewer comparisons, the more efficient the task is. So, when you have several select/omit comparisons, try to specify the one that selects or omits the most records first.

In the following examples, few records exist for which the *Rep* field is *JSMITH*. The examples show how to use DDS to select all the records before 1988 for a sales representative named *JSMITH* in the state of New York. All give the same results with different efficiency (in this example,

**3** is the most efficient).

```
|
| |...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
| |   A           S ST                          CMP(EQ 'NY') 1
| |   A           REP                          CMP(EQ 'JSMITH')
| |   A           YEAR                         CMP(LT 88)
| |   A
|
```

```
|
| |...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
| |   A           O YEAR                       CMP(LE 88) 2
| |   A           S ST                          CMP(EQ 'NY')
| |   A           REP                          CMP(EQ 'JSMITH')
| |   A
|
```

```
|
| |...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
| |   A           O REP                       CMP(NE 'JSMITH') 3
| |   A           O ST                         CMP(NE 'NY')
| |   A           S YEAR                       CMP(LT 88)
| |   A
|
```

Figure 3-3. Three Ways to Code Select/Omit Function

- 1** All records must be compared with all of the select fields *St*, *Rep*, and *Year* before they can be selected or omitted.
- 2** All records are compared with the *Year* field. Then, the records before 1988 have to be compared with the *St* and *Rep* fields.
- 3** All records are compared with the *Rep* field. Then, only the few for JSMITH are compared with the *St* field. Then, the few records that are left are compared to the *Year* field.

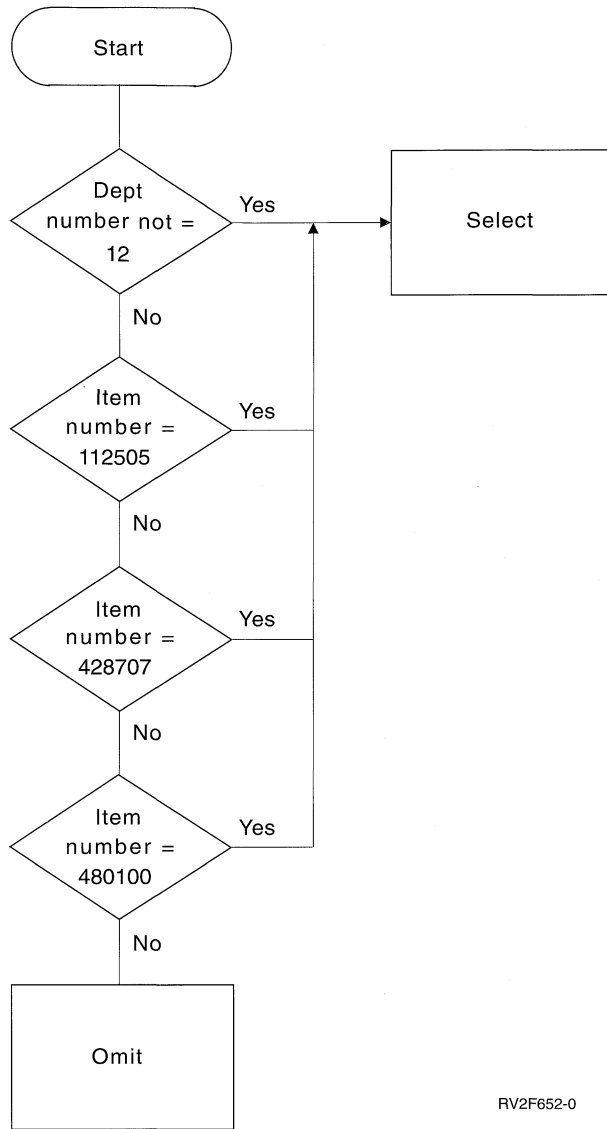
As another example, assume that you want to select the following:

- All records for departments other than Department 12.
- Only those records for Department 12 that contain an item number 112505, 428707, or 480100. No other records for Department 12 are to be selected.

|  
|  
|  
|  
|

If you create the preceding example with a sort sequence table, the select/omit fields are translated according to the sort table before the comparison. For example, with a sort sequence table using shared weightings for uppercase and lowercase, NY and ny are equal. For details, see the *DDS Reference*

The following diagram shows the logic included in this example:



RV2F652-0

The following shows how to code this example using the DDS select and omit functions:

```

|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
|   A       S DPTNBR          CMP(NE 12)
|   A       S ITMNBR          VALUES(112505 428707 480100)
|   A
  
```

It is possible to have an access path with select/omit values and process the file in arrival sequence. For example, a high-level language program can specify that the keyed access path is to be ignored. In this case, every record is read from the file in arrival sequence, but only those records meeting the select/omit values specified in the file are returned to the high-level language program.

A logical file with key fields and select/omit values specified can be processed in arrival sequence or using relative record numbers randomly. Records omitted by the select/omit values are not processed. That is, if an omitted record is requested by relative record number, the record is not returned to the high-level language program.

The system does not ensure that any additions or changes through a logical file will allow the record to be accessed again in the same logical file. For example, if the selection values of the logical file specifies only records with an A in *Fld1* and the program updates the record with a B in *Fld1*, the program cannot retrieve the record again using this logical file.

**Note:** You cannot select or omit based on the values of a floating-point field.

The two kinds of select/omit operations are: access path select/omit and dynamic select/omit. The default is access path select/omit. The select/omit specifications themselves are the same in each kind, but the system actually does the work of selecting and omitting records at different times.

### **Access Path Select/Omit**

With access path select/omit, the access path only contains keys that meet the select/omit values specified for the logical file. When you specify key fields for a file, an access path is kept for the file and maintained by the system when you add or update records in the physical file(s) used by the logical file. The only index entries in the access path are those that meet the select/omit values.

### **Dynamic Select/Omit**

With dynamic select/omit, when a program reads records from the file, the system only returns those records that meet the select/omit values. That is, the actual select/omit processing is done when records are read by a program, rather than when the records are added or changed. However, the keyed sequence access path contains all the keys, not just keys from selected records. Access paths using dynamic select/omit allow more access path sharing, which can improve performance. For more information about access path sharing, see “Using Existing Access Paths” on page 3-13.

To specify dynamic select/omit, use the dynamic selection (DYNSLT) keyword. With dynamic select/omit, key fields are not required.

If you have a file that is updated frequently and read infrequently, you may not need to update the access path for select/omit purposes until your program reads the file. In this case, dynamic select/omit might be the correct choice. The following example helps describe this.

You use a code field (A=active, I=inactive), which is changed infrequently, to select/omit records. Your program processes the active records and the majority (over 80%) of the records are active. It can be more efficient to use DYNSLT to dynamically select records at processing time rather than perform access path maintenance when the code field is changed.

### **Using the Open Query File Command to Select/Omit Records**

Another method of selecting records is using the QRYSLT parameter on the Open Query File (OPNQRYF) command. The open data path created by the OPNQRYF command is like a temporary logical file; that is, it is automatically deleted when it is closed. A logical file, on the other hand, remains in existence until you specifically delete it. For more details about the OPNQRYF command, see “Using the Open Query File (OPNQRYF) Command” on page 6-3.



## Using Existing Access Paths

When two or more files are based on the same physical files and the same key fields in the same order, they automatically share the same keyed sequence access path. When access paths are shared, the amount of system activity required to maintain access paths and the amount of auxiliary storage used by the files is reduced.

When a logical file with a keyed sequence access path is created, the system always tries to share an existing access path. For access path sharing to occur, an access path must exist on the system that satisfies the following conditions:

- The logical file member to be added must be based on the same physical file members that the existing access path is based on.
- The length, data type, and number of decimal positions specified for each key field must be identical in both the new file and the existing file.
- If the FIFO, LIFO, or FCFO keyword is not specified, the new file can have fewer key fields than the existing access paths. That is, a new logical file can share an existing access path if the beginning part of the key is identical. However, when a file shares a partial set of keys from an existing access path, any record updates made to fields that are part of the set of keys for the shared access path may change the record position in that access path. See “Example of Implicitly Shared Access Paths” on page 3-14 for a description of such a circumstance.
- The attributes of the access path (such as UNIQUE, LIFO, FIFO, or FCFO) and the attributes of the key fields (such as DESCEND, ABSVAL, UNSIGNED, and SIGNED) must be identical.

Exceptions:

1. A FIFO access path can share an access path in which the UNIQUE keyword is specified if all the other requirements for access path sharing are met.
  2. A UNIQUE access path can share a FIFO access path that needs to be rebuilt (for example, has \*REBLD maintenance specified), if all the other requirements for access path sharing are met.
- If the new logical file has select/omit specifications, they must be identical to the select/omit specifications of the existing access path. However, if the new logical file specifies DYNSTL, it can share an existing access path if the existing access path has either:
    - The dynamic select (DYNSTL) keyword specified
    - No select/omit keywords specified
  - The alternative collating sequence (ALTSEQ keyword) and the translation table (TRNTBL keyword) of the new logical file member, if any, must be identical to the alternative collating sequence and translation table of the existing access path.

**Note:** Logical files that contain concatenated or substring fields cannot share access paths with physical files.

The owner of any access path is the logical file member that originally created the access path. For a shared access path, if the logical member owning the access path is deleted, the first member to share the access path becomes the new owner. The FRCACPTH, MAINT, RECOVER, and UNIT parameters on the Create Logical File (CRTLF) command need not match the same parameters on an existing access path for that access path to be shared. When an access path is shared by several logical file members, and the FRCACPTH, MAINT, RECOVER,

and UNIT parameters are not identical, the system maintains the access path by the most restrictive value for each of the parameters specified by the sharing members. The following illustrates how this occurs:

MBRA specifies the following:    MBRB specifies the following:    System does the following:

FRCACCPH (*NO) MAINT (*IMMED) RECOVER (*AFTIPL)	FRCACCPH (*YES) MAINT (*DLY) RECOVER (*NO)	FRCACCPH (*YES) MAINT (*IMMED) RECOVER (*AFTIPL)
-------------------------------------------------------	--------------------------------------------------	--------------------------------------------------------

RSLH222-3

The UNIT parameter for the access path is always taken from the owning logical member.

Access path sharing does not depend on sharing between members; therefore, it does not restrict the order in which members can be deleted.

The Display File Description (DSPFD) and Display Database Relations (DSPDBR) commands show access path sharing relationships.

### Example of Implicitly Shared Access Paths

The purpose of this example is help you fully understand implicit access path sharing.

Two logical files, LFILE1 and LFILE2, are built over the physical file PFILE. LFILE1, which was created first, has two key fields, KFD1 and KFD2. LFILE2 has three key fields, KFD1, KFD2, and KFD3. The two logical files use two of the same key fields, but no access path is shared because the logical file with three key fields was created after the file with two key fields.

Figure 3-4. Physical and Logical Files Before Save and Restore

	Physical File (PFILE)	Logical File 1 (LFILE1)	Logical File 2 (LFILE2)
<b>Access Path</b>		KFD1, KFD2	KFD1, KFD2, KFD3
<b>Fields</b>	KFD1, KFD2, KFD3, A, B, C, D, E, F, G	KFD1, KFD2, KFD3, F, C, A	KFD1, KFD2, KFD3, D, G, F, E

An application uses LFILE1 to access the records and to change the KFD3 field to blank if it contains a C, and to a C if it is blank. This application causes the user no unexpected results because the access paths are not shared. However, after a save and restore of the physical file and both logical files, the program appears to do nothing and takes longer to process.

Unless you do something to change the restoration, the AS/400 system:

- Restores the logical file with the largest number of keys first
- Does not build unnecessary access paths

See "Controlling When Access Paths Are Rebuilt" on page 13-6 for information on changing these conditions.

Because it has three key fields, LFILE2 is restored first. After recovery, LFILE1 implicitly shares the access path for LFILE2. Users who do not understand implicitly shared access paths do not realize that when they use LFILE1 after a recovery, they are really using the key for LFILE2.

Figure 3-5. Physical and Logical Files After Save and Restore. Note that the only difference from before the save and restore is that the logical files now share the same access path.

	Physical File (PFILE)	Logical File 1 (LFILE1)	Logical File 2 (LFILE2)
<b>Access Path</b>		KFD1, KFD2, KFD3	KFD1, KFD2, KFD3
<b>Fields</b>	KFD1, KFD2, KFD3, A, B, C, D, E, F, G	KFD1, KFD2, KFD3, F, C, A	KFD1, KFD2, KFD3, D, G, F, E

The records to be tested and changed contain:

Relative Record	KFD1	KFD2	KFD3
001	01	01	<blank>
002	01	01	<blank>
003	01	01	<blank>
004	01	01	<blank>

The first record is read via the first key of 0101<blank> and changed to 0101C. The records now look like:

Relative Record	KFD1	KFD2	KFD3
001	01	01	C
002	01	01	<blank>
003	01	01	<blank>
004	01	01	<blank>

When the application issues a get next key, the next higher key above 0101<blank> is 0101C. This is the record that was just changed. However, this time the application changes the KFD3 field from C to blank.

Because the user does not understand implicit access path sharing, the application accesses and changes every record twice. The end result is that the application takes longer to run, and the records look like they have not changed.

## Creating a Logical File

Before creating a logical file, the physical file or files on which the logical file is based must already exist.

To create a logical file, take the following steps:

1. Type the DDS for the logical file into a source file. This can be done using SEU or another method. See "Working with Source Files" on page 14-3, for



```

|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
A* ORDER HEADER FILE (ORDHDRP) - PHYSICAL FILE RECORD DEFINITION
A                                     REF(DSTREFP)
A          R ORDHDR                   TEXT('Order header record')
A          CUST                       R
A          ORDER                      R
A          ORDATE                     R
A          CUSORD                     R
A          SHPVIA                     R
A          ORDSTS                     R
A          OPRNME                     R
A          ORDMNT                     R
A          CUTYPE                     R
A          INVNBR                     R
A          PRDAT                      R
A          SEQNBR                     R
A          OPNSTS                     R
A          LINES                      R
A          ACTMTH                     R
A          ACTYR                      R
A          STATE                      R
A

```

Figure 3-7. DDS for a Physical File (ORDHDRP) Built from a Field Reference File

The following example shows how to create a logical file ORDFILL with two record formats. One record format is defined for order header records from the physical file ORDHDRP; the other is defined for order detail records from the physical file ORDDTLP. (Figure 3-6 on page 3-16 shows the DDS for the physical file ORDDTLP, Figure 3-7 shows the DDS for the physical file ORDHDRP, and Figure 3-8 shows the DDS for the logical file ORDFILL.)

The logical file record format ORDHDR uses one key field, *Order*, for sequencing; the logical file record format ORDDTL uses two keys fields, *Order* and *Line*, for sequencing.

```

|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
A* ORDER TRANSACTION LOGICAL FILE (ORDFILL)
A          R ORDHDR                   PFILE(ORDHDRP)
A          K ORDER
A
A          R ORDDTL                   PFILE(ORDDTLP)
A          K ORDER
A          K LINE
A

```

Figure 3-8. DDS for the Logical File ORDFILL

To create the logical file ORDFILL with two associated physical files, use a Create Logical File (CRTLF) command like the following:

```

CRTLF FILE(DSTPRODLB/ORDFILL)
      TEXT('Order transaction logical file')

```

The DDS source is in the member ORDFILL in the file QDDSSRC. The file ORDFILL with a member of the same name is placed in the DSTPRODLB library. The access path for the logical file member ORDFILL arranges records from both the ORDHDRP and ORDDTLP files. Record formats for both physical files are keyed on *Order* as the common field. Because of the order in which they were specified in the logical file description, they are merged in *Order* sequence with duplicates between files retrieved first from the header file ORDHDRP and second

from the detail file ORDDTLP. Because FIFO, LIFO, or FCFO are not specified, the order of retrieval of duplicate keys in the same file is not guaranteed.

**Note:** In certain circumstances, it is better to use multiple logical files, rather than to use a multiple-format logical file. For example, when keyed access is used with a multiple-format logical file, it is possible to experience poor performance if one of the files has very few records. Even though there are multiple formats, the logical file has only one index, with entries from each physical file. Depending on the kind of processing being done by the application program (for example, using RPG SETLL and READE with a key to process the small file), the system might have to search all index entries in order to find an entry from the small file. If the index has many entries, searching the index might take a long time, depending on the number of keys from each file and the sequence of keys in the index. (If the small file has no records, performance is not affected, because the system can take a fast path and avoid searching the index.)

### Controlling How Records Are Retrieved in a File with Multiple Formats

In a logical file with more than one record format, key field definitions are required. Each record format has its own key definition, and the record format key fields can be defined to merge the records of the different formats. Each record format does not have to contain every key field in the key. Consider the following records:

#### Header Record Format:

Record	Order	Cust	Ordate
1	41882	41394	050688
2	32133	28674	060288

#### Detail Record Format:

Record	Order	Line	Item	Qtyord	Extens
A	32133	01	46412	25	125000
B	32133	03	12481	4	001000
C	41882	02	46412	10	050000
D	32133	02	14201	110	454500
E	41882	01	08265	40	008000

In DDS, the header record format is defined before the detail record format. If the access path uses the *Order* field as the first key field for both record formats and the *Line* field as the second key field for only the second record format, both in ascending sequence, the order of the records in the access path is:

Record 2  
 Record A  
 Record D  
 Record B  
 Record 1  
 Record E  
 Record C

**Note:** Records with duplicate key values are arranged first in the sequence in which the physical files are specified. Then, if duplicates still exist within a

record format, the duplicate records are arranged in the order specified by the FIFO, LIFO, or FCFO keyword. For example, if the logical file specified the DDS keyword FIFO, then duplicate records within the format would be presented in first-in-first-out sequence.

For logical files with more than one record format, you can use the \*NONE DDS function for key fields to separate records of one record format from records of other record formats in the same access path. Generally, records from all record formats are merged based on key values. However, if \*NONE is specified in DDS for a key field, only the records with key fields that appear in all record formats before the \*NONE are merged.

The logical file in the following example contains three record formats, each associated with a different physical file:

Record Format	Physical File	Key Fields
EMPMSTR	EMPMSTR	Empnbr (employee number) <b>1</b>
EMPHIST	EMPHIST	Empnbr, Empdat (employed date) <b>2</b>
EMPEDUC	EMPEDUC	Empnbr, Clsnbr (class number) <b>3</b>

**Note:** All record formats have one key field in common, the *Empnbr* field.

The DDS for this example is:

```

| ...+....1.....2.....3.....4.....5.....6.....7.....8
| A
| A      K EMPNBR   1
| A
| A      K EMPNBR   2
| A      K EMPDAT
| A
| A      K EMPNBR   3
| A      K *NONE
| A      K CLSNBR
| A

```

\*NONE is assumed for the second and third key fields for EMPMSTR and the third key field for EMPHIST because no key fields follow these key field positions.

The following shows the arrangement of the records:

Empnbr	Empdat	Clsnbr	Record Format Name
426			EMPMSTR
426	6/15/74		EMPHIST
426		412	EMPEDUC
426		520	EMPEDUC
427			EMPMSTR
427	9/30/75		EMPHIST
427		412	EMPEDUC

\*NONE serves as a separator for the record formats EMPHIST and EMPEDUC. All the records for EMPHIST with the same *Empnbr* field are grouped together and

sorted by the *Empdat* field. All the records for EMPEDUC with the same *Empnbr* field are grouped together and sorted by the *Clsnbr* field.

**Note:** Because additional key field values are placed in the key sequence access path to guarantee the above sequencing, duplicate key values are not predictable.

See the *DDS Reference* for additional examples of the \*NONE DDS function.

### **Controlling How Records Are Added to a File with Multiple Formats**

To add a record to a multiple format logical file, identify the member of the based-on physical file to which you want the record written. If the application you are using does not allow you to specify a particular member within a format, each of the formats in the logical file needs to be associated with a single physical file member. If one or more of the based-on physical files contains more than one member, you need to use the DTAMBRS parameter, described in “Logical File Members,” to associate a single member with each format. Finally, give each format in the multiple format logical file a unique name. If the multiple format logical file is defined in this way, then when you specify a format name on the add operation, you target a particular physical file member into which the record is added.

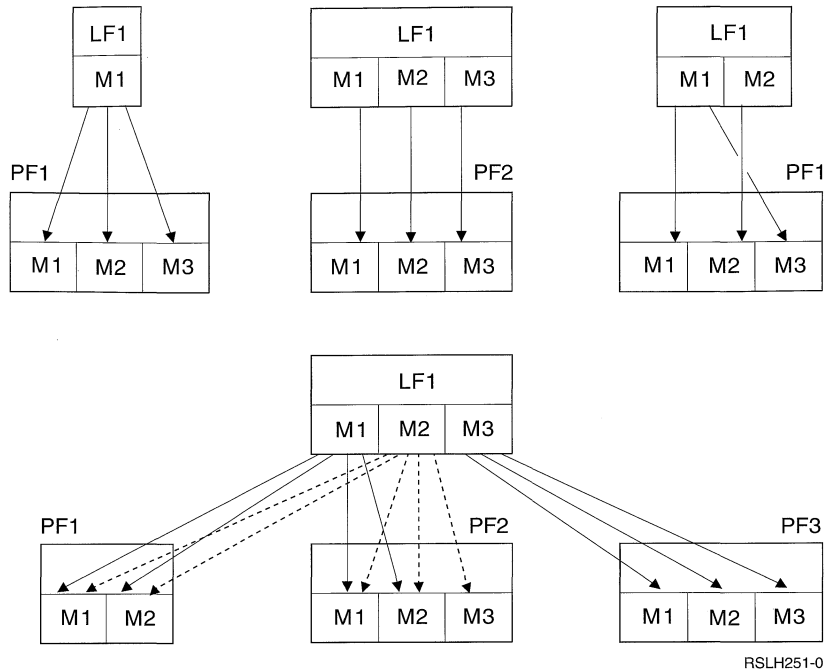
When you add records to a multiple-format logical file and your application program uses a file name instead of a record format name, you need to write a format selector program. For more information about format selector programs, see “Identifying Which Record Format to Add in a File with Multiple Formats” on page 7-9.

## **Logical File Members**

You can define members in logical files to separate the data into logical groups. The logical file member can be associated with one physical file member or with several physical file members.

The following illustrates this concept:





The record formats used with all logical members in a logical file must be defined in DDS when the file is created. If new record formats are needed, another logical file or record format must be created.

The attributes of an access path are determined by information specified in DDS and on commands when the logical file is created. The selection of data members is specified in the DTAMBRS parameter on the Create Logical File (CRTLF) and Add Logical File Member (ADDLFM) commands.

When a logical file is defined, the physical files used by the logical file are specified in DDS by the record level PFILE or JFILE keyword. If multiple record formats are defined in DDS, a PFILE keyword must be specified for each record format. You can specify one or more physical files for each PFILE keyword.

When a logical file is created or a member is added to the file, you can use the DTAMBRS parameter on the Create Logical File (CRTLF) or the Add Logical File Member (ADDLFM) command to specify which members of the physical files used by the logical file are to be used for data. \*NONE can be specified as the physical file member name if no members from a physical file are to be used for data.



more members, you must specify more than 1 for the MAXMBRS parameter on the CRTLF command. The following example of adding a member to a logical file uses the CRTLF command used earlier in “Creating a Logical File” on page 3-15.

```
CRTLF FILE(DSTPRODLB/ORDHDRL)
      MBR(*FILE) DTAMBR(*ALL)
      TEXT('Order header logical file')
```

\*FILE is the default for the MBR parameter and means the name of the member is the same as the name of the file. All the members of the associated physical file (ORDHDRP) are used in the logical file (ORDHDRL) member. The text description is the text description of the member.

---

## Join Logical File Considerations

This section covers the following topics:

- Basic concepts of joining two physical files (Example 1)
- Setting up a join logical file
- Using more than one field to join files (Example 2)
- Handling duplicate records in secondary files using the JDUPSEQ keyword (Example 3)
- Handling join fields whose attributes do not match (Example 4)
- Using fields that never appear in the record format to join files—neither fields (Example 5)
- Specifying key fields in join logical files (Example 6)
- Specifying select/omit statements in join logical files
- Joining three or more physical files (Example 7)
- Joining a physical file to itself (Example 8)
- Using default data for records missing from secondary files—the JDFTVAL keyword (Example 9)
- Describing a complex join logical file (Example 10)
- Performance considerations
- Data integrity considerations
- Summary of rules for join logical files


In general, the examples in this section include a picture of the files, DDS for the files, and sample data. For Example 1, several cases are given that show how to join files in different situations (when data in the physical files varies).

In the examples, for convenience and ease of recognition, join logical files are shown with the label JLF, and physical files are illustrated with the labels PF1, PF2, PF3, and so forth.

### Basic Concepts of Joining Two Physical Files (Example 1)

A **join logical file** is a logical file that combines (in one record format) fields from two or more physical files. In the record format, not all the fields need to exist in all the physical files.

The following example illustrates a join logical file that joins two physical files. This example is used for the five cases discussed in Example 1.



In this example, employee number is common to both physical files (PF1 and PF2), but name is found only in PF1, and salary is found only in PF2.

With a join logical file, the application program does one read operation (to the record format in the join logical file) and gets all the data needed from both physical files. Without the join specification, the logical file would contain two record formats, one based on PF1 and the other based on PF2, and the application program would have to do two read operations to get all the needed data from the two physical files. Thus, join provides more flexibility in designing your database.

However, a few restrictions are placed on join logical files:

- You cannot change a physical file through a join logical file. To do update, delete, or write (add) operations, you must create a second multiple format logical file and use it to change the physical files. You can also use the physical files, directly, to do the change operations.
- You cannot use DFU to display a join logical file.
- You can specify only one record format in a join logical file.
- The record format in a join logical file cannot be shared.
- A join logical file cannot share the record format of another file.
- Key fields must be fields defined in the join record format and must be fields from the first file specified on the JFILE keyword (which is called the primary file).
- Select/omit fields must be fields defined in the join record format, but can come from any of the physical files.
- Commitment control cannot be used with join logical files.

The following shows the DDS for Example 1:

```

JLF
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R JOINREC          JFILE(PF1 PF2)
  A          J                  JOIN(PF1 PF2)
  A          JFLD(NBR NBR)
  A          NBR                 JREF(PF1)
  A          NAME
  A          SALARY
  A          K NBR
  A

PF1
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R REC1
  A          NBR                10
  A          NAME                20
  A          K NBR
  A

PF2
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R REC2
  A          NBR                10
  A          SALARY              7 2
  A          K NBR
  A

```

Figure 3-9. DDS Example for Joining Two Physical Files

The following describes the DDS for the join logical file in Example 1 (see the *DDS Reference* for more information on the specific keywords):

The record level specification identifies the record format name used in the join logical file.

**R** Identifies the record format. Only one record format can be placed in a join logical file.

**JFILE** Replaces the **PF1** keyword used in simple and multiple-format logical files. You must specify at least two physical files. The first file specified on the **JFILE** keyword is the **primary file**. The other files specified on the **JFILE** keyword are **secondary files**.

The join specification describes the way a pair of physical files is joined. The second file of the pair is always a secondary file, and there must be one join specification for each secondary file.

**J** Identifies the start of a join specification. You must specify at least one join specification in a join logical file. A join specification ends at the first field name specified in positions 19 through 28 or at the next **J** specified in position 17.

**JOIN** Identifies which two files are joined by the join specification. If only two physical files are joined by the join logical file, the **JOIN** keyword is optional. See “Joining Three or More Physical Files (Example 7)” on page 3-40 later in this section for an example of how to use this keyword.

**JFLD** Identifies the **join fields** that join records from the physical files specified on the JOIN. JFLD must be specified at least once for each join specification. The join fields are fields common to the physical files. The first join field is a field from the first file specified on the JOIN keyword, and the second join field is a field from the second file specified on the JOIN keyword.

Join fields, except character type fields, must have the same attributes (data type, length, and decimal positions). If the fields are character type fields, they do not need to have the same length. If you are joining physical file fields that do not have the same attributes, you can redefine them for use in a join logical file. See “Using Join Fields Whose Attributes Are Different (Example 4)” on page 3-36 for a description and example.

The field level specification identifies the fields included in the join logical file.

**Field names** Specifies which fields (in this example, *Nbr*, *Name*, and *Salary*) are used by the application program. At least one field name is required. You can specify any field names from the physical files used by the logical file. You can also use keywords like RENAME, CONCAT, or SST as you would in simple and multiple format logical files.

**JREF** In the record format (which follows the join specification level and precedes the key field level, if any), the field names must uniquely identify which physical file the field comes from. In this example, the *Nbr* field occurs in both PF1 and PF2. Therefore, the JREF keyword is required to identify the file from which the *Nbr* field description will be used.

The key field level specification is optional, and includes the key field names for the join logical file.

**K** Identifies a key field specification. The K appears in position 17. Key field specifications are optional.

**Key field names**

Key field names (in this example, *Nbr* is the only key field) are optional and make the join logical file an indexed (keyed sequence) file. Without key fields, the join logical file is an arrival sequence file. In join logical files, key fields must be fields from the primary file, and the key field name must be specified in positions 19 through 28 in the logical file record format.

The select/omit field level specification is optional, and includes select/omit field names for the join logical file.

**S or O** Identifies a select or omit specification. The S or O appears in position 17. Select/omit specifications are optional.

**Select/omit field names**

Only those records meeting the select/omit values will be returned to the program using the logical file. Select/omit fields must be specified in positions 19 through 28 in the logical file record format.

## Reading a Join Logical File

The following cases describe how the join logical file in Figure 3-9 on page 3-25 presents records to an application program.

The PF1 file is specified first on the JFILE keyword, and is therefore the primary file. When the application program requests a record, the system does the following:

1. Uses the value of the first join field in the primary file (the *Nbr* field in PF1).
2. Finds the first record in the secondary file with a matching join field (the *Nbr* field in PF2 matches the *Nbr* field in PF1).
3. For each match, joins the fields from the physical files into one record and provides this record to your program. Depending on how many records are in the physical files, one of the following conditions could occur:
  - a. For all records in the primary file, only one matching record is found in the secondary file. The resulting join logical file contains a single record for each record in the primary file. See “Matching Records in Primary and Secondary Files (Case 1)” on page 3-28.
  - b. For some records in the primary file, no matching record is found in the secondary file.

If you specify the JDFTVAL keyword:

- For those records in the primary file that have a matching record in the secondary file, the system joins to the secondary, or multiple secondaries. The result is one or more records for each record in the primary file.
- For those records in the primary file that do not have a matching record in the secondary file, the system adds the default value fields for the secondary file and continues the join operation. You can use the DFT keyword in the physical file to define which defaults are used. See “Record Missing in Secondary File; JDFTVAL Keyword Not Specified (Case 2A)” on page 3-28 and “Record Missing in Secondary File; JDFTVAL Keyword Specified (Case 2B)” on page 3-29.

**Note:** If the DFT keyword is specified in the secondary file, the value specified for the DFT keyword is used in the join. The result would be at least one join record for each primary record.

- If a record exists in the secondary file, but the primary file has no matching value, no record is returned to your program. A second join logical file can be used that reverses the order of primary and secondary files to determine if secondary file records exist with no matching primary file records.

If you do not specify the JDFTVAL keyword:

- If a matching record in a secondary file exists, the system joins to the secondary, or multiple secondaries. The result is one or more records for each record in the primary file.
- If a matching record in a secondary file does not exist, the system does not return a record.

**Note:** When the JDFTVAL is not specified, the system returns a record only if a match is found in every secondary file for a record in the primary file.

In the following examples, cases 1 through 4 describe sequential read operations, and case 5 describes reading by key.

### Matching Records in Primary and Secondary Files (Case 1)

Assume that a join logical file is specified as in Figure 3-9 on page 3-25, and that four records are contained in both PF1 and PF2, as follows:

PF1		PF2	
235	Anne	235	1700.00
440	Doug	440	950.50
500	Mark	500	2100.00
729	Sue	729	1400.90

RSLH255-0

The program does four read operations and gets the following records:

JLF		
235	Anne	1700.00
440	Doug	950.50
500	Mark	2100.00
729	Sue	1400.90

RSLH256-0

### Record Missing in Secondary File; JDFTVAL Keyword Not Specified (Case 2A)

Assume that a join logical file is specified as in Figure 3-9 on page 3-25, and that there are four records in PF1 and three records in PF2, as follows:

PF1		PF2	
235	Anne	235	1700.00
440	Doug	440	950.50
500	Mark		
729	Sue	729	1400.90

No record was found for number 500 in PF2.

RSLH257-1

With the join logical file shown in Example 1, the program reads the join logical file and gets the following records:

JLF		
235	Anne	1700.00
440	Doug	950.50
729	Sue	1400.90

RSLH258-0

If you do not specify the JDFTVAL keyword and no match is found for the join field in the secondary file, the record is not included in the join logical file.



## Record Missing in Secondary File; JDFTVAL Keyword Specified (Case 2B)

Assume that a join logical file is specified as in Figure 3-9 on page 3-25, except that the JDFTVAL keyword is specified, as shown in the following DDS:

```

JLF
|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
  A                               JDFTVAL
  A      R JOINREC                JFILE(PF1 PF2)
  A      J                        JOIN(PF1 PF2)
  A                               JFLD(NBR NBR)
  A      NBR                      JREF(PF1)
  A      NAME
  A      SALARY
  A      K NBR
  A
  
```

The program reads the join logical file and gets the following records:

```

JLF
235  Anne  1700.00
440  Doug   950.50
500  Mark  0000.00
729  Sue   1400.90
  
```

← A record for number 500 is returned if JDFTVAL is specified (but the SALARY is 0).

RSLH260-0

With JDFTVAL specified, the system returns a record for 500, even though the record is missing in PF2. Without that record, some field values can be missing in the join record (in this case, the *Salary* field is missing). With JDFTVAL specified, missing character fields normally use blanks; missing numeric fields use zeros. However, if the DFT keyword is specified for the field in the physical file, the default value specified on the DFT keyword is used.

## Secondary File Has More Than One Match for a Record in the Primary File (Case 3)

Assume that a join logical file is specified as in Figure 3-9 on page 3-25, and that four records in PF1 and five records in PF2, as follows:

PF1	PF2
235 Anne	235 1700.00
440 Doug	235 1500.00
500 Mark	440 950.50
729 Sue	500 2100.00
	729 1400.90

← Duplicate record was found in PF2 for number 235.

RSLH261-1

The program gets five records:

JLF

235	Anne	1700.00
235	Anne	1500.00
440	Doug	950.50
500	Mark	2100.00
729	Sue	1400.90

Order of records received for 235 is unpredictable unless you specify the JDUPSEQ keyword.

RSLH262-0

For more information, see "Reading Duplicate Records in Secondary Files (Example 3)" on page 3-34.

### Extra Record in Secondary File (Case 4)

Assume that a join logical file is specified as in Figure 3-9 on page 3-25, and that four records are contained in PF1 and five records in PF2, as follows:

PF1

235	Anne
440	Doug
500	Mark
729	Sue

PF2

235	1700.00
301	1500.00
440	950.50
500	2100.00
729	1400.90

Record for number 301 is only in PF2.

RSLH263-1

The program reads the join logical file and gets only four records, which would be the same even if JDFTVAL was specified (because a record must always be contained in the primary file to get a join record):

JLF

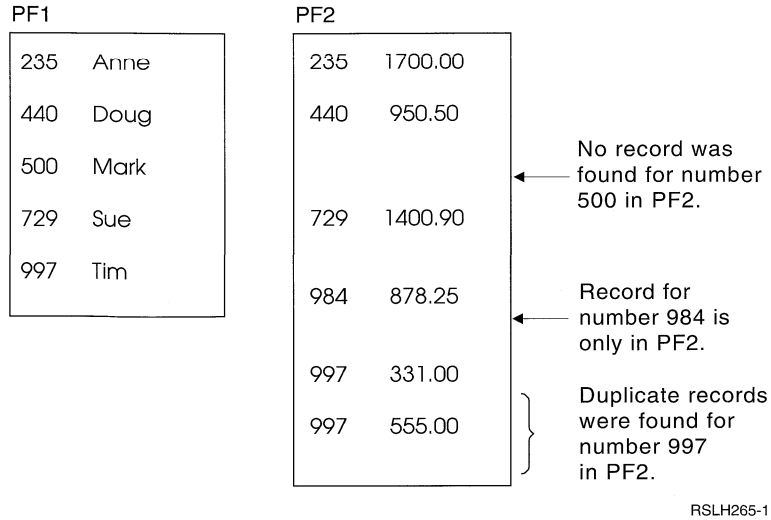
235	Anne	1700.00
440	Doug	950.50
500	Mark	2100.00
729	Sue	1400.90

RSLH264-0

### Random Access (Case 5)

Assume that a join logical file is specified as in Figure 3-9 on page 3-25. Note that the join logical file has key fields defined. This case shows which records would be returned for a random access read operation using the join logical file.

Assume that PF1 and PF2 have the following records:



The program can get the following records:

Given a value of 235 from the program for the *Nbr* field in the logical file, the system supplies the following record:

235	Anne	1700.00
-----	------	---------

RSLH266-0

Given a value of 500 from the program for the *Nbr* field in the logical file and with the JDFTVAL keyword specified, the system supplies the following record:

500	Mark	0.00
-----	------	------

RSLH267-0

**Note:** If the JDFTVAL keyword was not specified in the join logical file, no record would be found for a value of 500 because no matching record is contained in the secondary file.

Given a value of 984 from the program for the *Nbr* field in the logical file, the system supplies no record and a no record found exception occurs because record 984 is not in the primary file.

Given a value of 997 from the program for the *Nbr* field in the logical file, the system returns one of the following records:

997	Tim	331.00
-----	-----	--------

or

997	Tim	555.00
-----	-----	--------

RSLH268-0

Which record is returned to the program cannot be predicted. To specify which record is returned, specify the JDUPSEQ keyword in the join logical file. See "Reading Duplicate Records in Secondary Files (Example 3)" on page 3-34.

**Notes:**

1. With random access, the application programmer must be aware that duplicate records could be contained in PF2, and ensure that the program does more than one read operation for records with duplicate keys. If the program were using sequential access, a second read operation would get the second record.
2. If you specify the JDUPSEQ keyword, the system can create a separate access path for the join logical file (because there is less of a chance the system will find an existing access path that it can share). If you omit the JDUPSEQ keyword, the system can share the access path of another file. (In this case, the system would share the access path of PF2.)

## Setting Up a Join Logical File

To set up a join logical file, do the following:

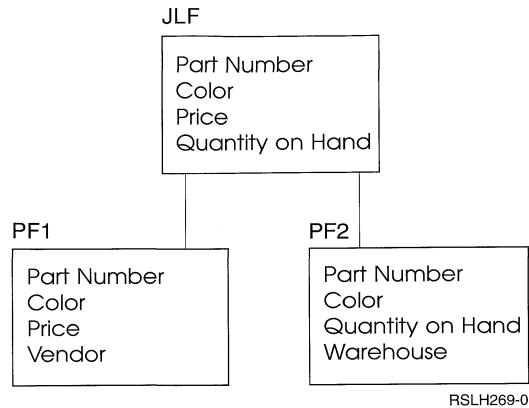
1. Find the field names of all the physical file fields you want in the logical file record format. (You can display the fields contained in files using the Display File Field Description [DSPFFD] command.)
2. Describe the fields in the record format. Write the field names in a vertical list. This is the start of the record format for the join logical file.

**Note:** You can specify the field names in any order. If the same field names appear in different physical files, specify the name of the physical file on the JREF keyword for those fields. You can rename fields using the RENAME keyword, and concatenate fields from the same physical file using the CONCAT keyword. A subset of an existing character, hexadecimal, or zoned decimal field can be defined using the SST keyword. The substring of a character or zoned decimal field is a character field, and the substring of a hexadecimal field is also a hexadecimal field. You can redefine fields: changing their data type, length, or decimal positions.

3. Specify the names of the physical files as parameter values on the JFILE keyword. The first name you specify is the primary file. The others are all secondary files. For best performance, specify the secondary files with the least records first after the primary file.
4. For each secondary file, code a join specification. On each join specification, identify which pair of files are joined (using the JOIN keyword; optional if only one secondary file), and identify which fields are used to join the pair (using the JFLD keyword; at least one required in each join specification).
5. Optionally, specify the following:
  - a. The JDFTVAL keyword. Do this if you want to return a record for each record in the primary file even if no matching record exists in a secondary file.
  - b. The JDUPSEQ keyword. Do this for fields that might have duplicate values in the secondary files. JDUPSEQ specifies on which field (other than one of the join fields) to sort these duplicates, and the sequence that should be used.
  - c. Key fields. Key fields cannot come from a secondary file. If you omit key fields, records are returned in arrival sequence as they appear in the primary file.
  - d. Select/omit fields. In some situations, you must also specify the dynamic selection (DYNSLT) keyword at the file level.
  - e. Neither fields. For a description, see “Describing Fields That Never Appear in the Record Format (Example 5)” on page 3-37.

## Using More Than One Field to Join Files (Example 2)

You can specify more than one join field to join a pair of files. The following shows the fields in the logical file and the two physical files.



The DDS for these files is as follows:

```

JLF
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R JOINREC          JFILE(PF1 PF2)
  A          J                   JOIN(PF1 PF2)
  A                                     JFLD(PTNBR PTNBR)
  A                                     JFLD(COLOR COLOR)
  A          PTNBR                JREF(PF1)
  A          COLOR                JREF(PF1)
  A          PRICE
  A          QUANTOH
  A
  
```

```

PF1
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R REC1
  A          PTNBR                4
  A          COLOR                20
  A          PRICE                7 2
  A          VENDOR                40
  A
  
```

```

PF2
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R REC2
  A          PTNBR                4
  A          COLOR                20
  A          QUANTOH              5 0
  A          WAREHSE              30
  A
  
```

Assume that the physical files have the following records:

PF1

100	Black	22.50	ABC Corp.
100	White	20.00	Ajax Inc.
120	Yellow	3.75	ABC Corp.
187	Green	110.95	ABC Corp.
187	Red	110.50	ABC Corp.
190	Blue	40.00	Ajax Inc.

PF2

100	Black	23	ABC Corp.
100	White	15	Ajax Inc.
120	Yellow	102	ABC Corp.
187	Green	0	ABC Corp.
187	Red	2	ABC Corp.
190	White	2	Ajax Inc.

RSLH271-0

If the file is processed sequentially, the program receives the following records:

JLF

100	Black	22.50	23
100	White	20.00	15
120	Yellow	3.75	102
187	Green	110.95	0
187	Red	110.50	2

RSLH272-0

Note that no record for part number 190, color blue, is available to the program, because a match was not found on both fields in the secondary file. Because JDFTVAL was not specified, no record is returned.

### Reading Duplicate Records in Secondary Files (Example 3)

Sometimes a join to a secondary file produces more than one record from the secondary file. When this occurs, specifying the JDUPSEQ keyword in the join specification for that secondary file tells the system to base the order of the duplicate records on the specified field in the secondary file.

The DDS for the physical files and for the join logical file are as follows:

```

JLF
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R JREC                      JFILE(PF1 PF2)
  A          J                          JOIN(PF1 PF2)
  A                                     JFLD(NAME1 NAME2)
  A                                     JDUPSEQ(TELEPHONE)
  A          NAME1
  A          ADDR
  A          TELEPHONE
  A

PF1
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R REC1
  A          NAME1          10
  A          ADDR          20
  A

PF2
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R REC2
  A          NAME2          10
  A          TELEPHONE      8
  A

```

Figure 3-10. DDS Example Using the JDUPSEQ Keyword

The physical files have the following records:

PF1

Anne	120 1st St.
Doug	40 Pillsbury
Mark	2 Lakeside Dr.

PF2

Anne	555-1111
Anne	555-6666
Anne	555-2222
Doug	555-5555

RSLH273-0

The join logical file returns the following records:

JLF

Anne	120 1st St.	555-1111	} Anne's telephone numbers are in ascending order.
Anne	120 1st St.	555-2222	
Anne	120 1st St.	555-6666	
Doug	40 Pillsbury	555-5555	

RSLH274-1

The program reads all the records available for Anne, then Doug, then Mark. Anne has one address, but three telephone numbers. Therefore, there are three records returned for Anne.

The records for Anne sort in ascending sequence by telephone number because the JDUPSEQ keyword sorts in ascending sequence unless you specify \*DESCEND as the keyword parameter. The following example shows the use of \*DESCEND in DDS:

JLF

```

|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
A          R JREC          JFILE(PF1 PF2)
A          J              JOIN(PF1 PF2)
A                          JFLD(NAME1 NAME2)
A                          JDUPSEQ(TELEPHONE *DESCEND)
A          NAME1
A          ADDR
A          TELEPHONE
A

```

When you specify JDUPSEQ with \*DESCEND, the records are returned as follows:

JLF

Anne	120 1st St.	555-6666	} Anne's telephone numbers are in descending order.
Anne	120 1st St.	555-2222	
Anne	120 1st St.	555-1111	
Doug	40 Pillsbury	555-5555	

RSLH277-1

**Note:** The JDUPSEQ keyword applies only to the join specification in which it is specified. For an example showing the JDUPSEQ keyword in a join logical file with more than one join specification, see "A Complex Join Logical File (Example 10)" on page 3-45.

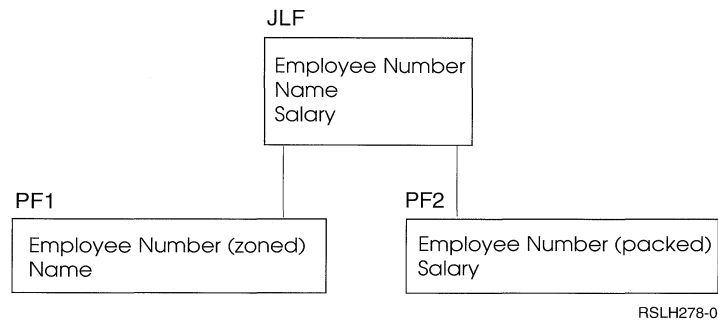
## Using Join Fields Whose Attributes Are Different (Example 4)

Fields from physical files that you are using as join fields generally have the same attributes (length, data type, and decimal positions). For example, as in Figure 3-10 on page 3-34, the *Name1* field is a character field 10 characters long in physical file PF1, and can be joined to the *Name2* field, a character field 10 characters long in physical file PF2. The *Name1* and *Name2* fields have the same characteristics and, therefore, can easily be used as join fields.

You can also use character type fields that have different lengths as join fields without requiring any redefinition of the fields. For example, if the NAME1 Field of PF1 was 10 characters long and the NAME2 field of PF2 was 15 characters long, those fields could be used as join fields without redefining one of the fields.

The following is an example in which the join fields do not have the same attributes. The *Nbr* field in physical file PF1 and the *Nbr* field in physical file PF2 both have a length of 3 specified in position 34, but in the PF1 file the field is zoned (S in position 35), and in the PF2 file the field is packed (P in position 35). To join the two files using these fields as join fields, you must redefine one or both fields to have the same attributes.

The following illustrates the fields in the logical and physical files:



The DDS for these files is as follows:



```

JLF
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R JOINREC          JFILE(PF1 PF2)
  A          J                   JOIN(PF1 PF2)
  A          J                   JFLD(NBR NBR)
  A          NBR                 S       JREF(2)
  A          NAME
  A          SALARY
  A

```

```

PF1
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R REC1
  A          NBR                 3S 0 <-Zoned
  A          NAME                 20
  A          K NBR
  A

```

```

PF2
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R REC2
  A          NBR                 3P 0 <-Packed
  A          SALARY              7 2
  A          K NBR
  A

```

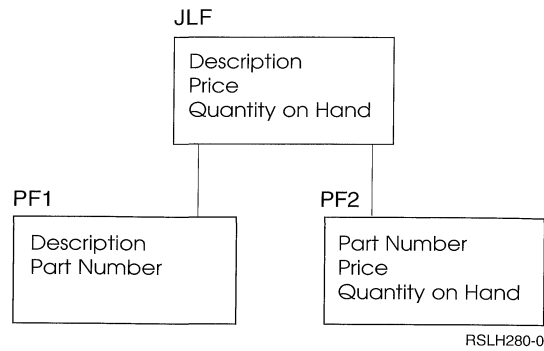
**Note:** In this example, the *Nbr* field in the logical file comes from PF2, because JREF(2) is specified. Instead of specifying the physical file name, you can specify a relative file number on the JREF keyword; in this example, the 2 indicates PF2.

Because the *Nbr* fields in the PF1 and PF2 files are used as the join fields, they must have the same attributes. In this example, they do not. Therefore, you must redefine one or both of them to have the same attributes. In this example, to resolve the difference in the attributes of the two employee number fields, the *Nbr* field in JLF (which is coming from the PF2 file) is redefined as zoned (S in position 35 of JLF).

## Describing Fields That Never Appear in the Record Format (Example 5)

A neither field (N specified in position 38) can be used in join logical files for neither input nor output. Programs using the join logical file cannot see or read neither fields. Neither fields are not included in the record format. Neither fields cannot be key fields or used in select/omit statements in the joined file. You can use a neither field for a join field (specified at the join specification level on the JFLD keyword) that is redefined at the record level only to allow the join, but is not needed or wanted in the program.

In the following example, the program reads the descriptions, prices, and quantity on hand of parts in stock. The part numbers themselves are not wanted except to bring together the records of the parts. However, because the part numbers have different attributes, at least one must be redefined.



The DDS for these files is as follows:

```

JLF
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R JOINREC                JFILE(PF1 PF2)
  A          J                        JOIN(PF1 PF2)
  A          JFLD(PRTNBR PRTNBR)
  A          PRTNBR          S N      JREF(1)
  A          DESC
  A          PRICE
  A          QUANT
  A          K DESC
  A
  
```

```

PF1
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R REC1
  A          DESC          30
  A          PRTNBR        6P 0
  A
  
```

```

PF2
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A          R REC2
  A          PRTNBR        6S 0
  A          PRICE          7 2
  A          QUANT          8 0
  A
  
```

In PF1, the *Prtnbr* field is a packed decimal field; in PF2, the *Prtnbr* field is a zoned decimal field. In the join logical file, they are used as join fields, and the *Prtnbr* field from PF1 is redefined to be a zoned decimal field by specifying an S in position 35 at the field level. The JREF keyword identifies which physical file the field comes from. However, the field is not included in the record format; therefore, N is specified in position 38 to make it a neither field. A program using this file would not see the field.

In this example, a sales clerk can type a description of a part. The program can read the join logical file for a match or a close match, and display one or more parts for the user to examine, including the description, price, and quantity. This application assumes that part numbers are not necessary to complete a customer order or to order more parts for the warehouse.

## Specifying Key Fields in Join Logical Files (Example 6)

If you specify key fields in a join logical file, the following rules apply:

- The key fields must exist in the primary physical file.
- The key fields must be named in the join record format in the logical file in positions 19 through 28.
- The key fields cannot be fields defined as neither fields (N specified in position 38 for the field) in the logical file.

The following illustrates the rules for key fields:

```

JLF
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
|
|  A          R JOINREC          JFILE(PF1 PF2)
|  A          J                  JOIN(PF1 PF2)
|  A          JFLD(NBR NUMBER)
|  A          JFLD(FLD3 FLD31)
|  A          FLD1                RENAME(F1)
|  A          FLD2                JREF(2)
|  A          FLD3                35  N
|  A          NAME
|  A          TELEPHONE          CONCAT(AREA LOCAL)
|  A          K FLD1
|  A          K NAME
|  A

```

```

PF1
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
|
|  A          R REC1
|  A          NBR                4
|  A          F1                20
|  A          FLD2              7  2
|  A          FLD3              40
|  A          NAME              20
|  A

```

```

PF2
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
|
|  A          R REC2
|  A          NUMBER            4
|  A          FLD2              7  2
|  A          FLD31            35
|  A          AREA              3
|  A          LOCAL             7
|  A

```

The following fields **cannot** be key fields:

- Nbr* (not named in positions 19 through 28)
- Number* (not named in positions 19 through 28)
- F1* (not named in positions 19 through 28)
- Fld31* (comes from a secondary file)
- Fld2* (comes from a secondary file)
- Fld3* (is a neither field)
- Area* and *Local* (not named in positions 19 through 28)
- Telephone* (is based on fields from a secondary file)

## Specifying Select/Omit Statements in Join Logical Files

If you specify select/omit statements in a join logical file, the following rules apply:

- The fields can come from any physical file the logical file uses (specified on the JFILE keyword).
- The fields you specify on the select/omit statements cannot be fields defined as neither fields (N specified in position 38 for the field).
- In some circumstances, you must specify the DYNSLT keyword when you specify select/omit statements in join logical files. For more information and examples, see the DYNSLT keyword in the *DDS Reference*.

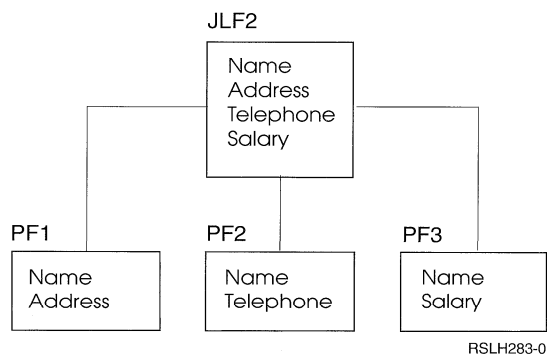
For an example showing select/omit statements in a join logical file, see “A Complex Join Logical File (Example 10)” on page 3-45.

## Joining Three or More Physical Files (Example 7)

You can use a join logical file to join as many as 32 physical files. These files must be specified on the JFILE keyword. The first file specified on the JFILE keyword is the primary file; the other files are all secondary files.

The physical files must be joined in pairs, with each pair described by a join specification. Each join specification must have one or more join fields identified.

The following shows the fields in the files and one field common to all the physical files in the logical file:



In this example, the *Name* field is common to all the physical files (PF1, PF2, and PF3), and serves as the join field.

The following shows the DDS for the physical and logical files:

```

JLF
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A      R JOINREC          JFILE(PF1 PF2 P3)
  A      J                  JOIN(PF1 PF2)
  A      J                  JFLD(NAME NAME)
  A      J                  JOIN(PF2 PF3)
  A      J                  JFLD(NAME NAME)
  A      NAME              JREF(PF1)
  A      ADDR
  A      TELEPHONE
  A      SALARY
  A      K NAME
  A

```

```

PF1
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A      R REC1
  A      NAME              10
  A      ADDR              20
  A      K NAME
  A

```

```

PF2
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A      R REC2
  A      NAME              10
  A      TELEPHONE        7
  A      K NAME
  A

```

```

PF3
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A      R REC3
  A      NAME              10
  A      SALARY            9 2
  A      K NAME
  A

```

Assume the physical files have the following records:

PF1

Anne	120 1st St.
Doug	40 Pillsbury
Mark	2 Lakeside Dr.
Tom	335 Elm St.

PF2

Anne	555-1111
Doug	555-5555
Mark	555-0000
Sue	555-3210

PF3

Anne	1700.00
Doug	950.00
Mark	2100.00

RSLH285-0

The program reads the following logical file records:

JLF

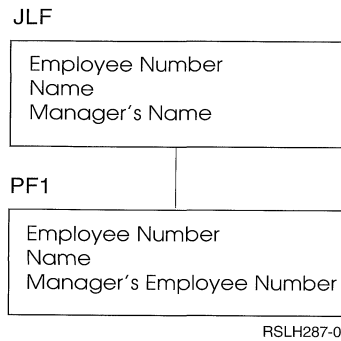
Anne	120 1st St.	555-1111	1700.00
Doug	40 Pillsbury	555-5555	950.00
Mark	2 Lakeside Dr.	555-0000	2100.00

RSLH286-0

No record is returned for Tom because a record is not found for him in PF2 and PF3 and the JDFTVAL keyword is not specified. No record is returned for Sue because the primary file has no record for Sue.

## Joining a Physical File to Itself (Example 8)

You can join a physical file to itself to read records that are formed by combining two or more records from the physical file itself. The following example shows how:



The following shows the DDS for these files:

```

JLF
|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
  A                               JDFTVAL
  A          R JOINREC            JFILE(PF1 PF1)
  A          J                    JOIN(1 2)
  A                               JFLD(MGRNBR NBR)
  A          NBR                   JREF(1)
  A          NAME                   JREF(1)
  A          MGRNAME                RENAME(NAME)
  A                               JREF(2)
  A

```

```

PF1
|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
  A          R RCD1
  A          NBR                    3
  A          NAME                    10      DFT('none')
  A          MGRNBR                  3
  A

```

### Notes:

1. Relative file numbers must be specified on the JOIN keyword because the same file name is specified twice on the JFILE keyword. Relative file number 1 refers to the first physical file specified on the JFILE keyword, 2 refers to the second, and so forth.
2. With the same physical files specified on the JFILE keyword, the JREF keyword is required for each field specified at the field level.

Assume the following records are contained in PF1:

PF1

235	Anne	440
440	Doug	729
500	Mark	440
729	Sue	888

RSLH289-0

The program reads the following logical file records:

JLF

235	Anne	Doug
440	Doug	Sue
500	Mark	Doug
729	Sue	none

RSLH290-0

Note that a record is returned for the manager name of Sue because the JDFTVAL keyword was specified. Also note that the value none is returned because the DFT keyword was used on the *Name* field in the PF1 physical file.

## Using Default Data for Missing Records from Secondary Files (Example 9)

If you are joining more than two files, and you specify the JDFTVAL keyword, the default value supplied by the system for a join field missing from a secondary file is used to join to other secondary files. If the DFT keyword is specified in the secondary file, the value specified for the DFT keyword is used in the logical file.

The DDS for the files is as follows:

```
JLF
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A                               JDFTVAL
  A                               JFILE(PF1 PF2 PF3)
  A           R JRCD              JOIN(PF1 PF2)
  A           J                   JFLD(NAME NAME)
  A           J                   JOIN(PF2 PF3)
  A                               JFLD(TELEPHONE TELEPHONE)
  A           NAME                JREF(PF1)
  A           ADDR
  A           TELEPHONE            JREF(PF2)
  A           LOC
  A
```

```
PF1
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A           R RCD1
  A           NAME                20
  A           ADDR                40
  A           COUNTRY             40
  A
```

```
PF2
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A           R RCD2
  A           NAME                20
  A           TELEPHONE            8           DFT('999-9999')
  A
```

```
PF3
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...8
  A           R RCD3
  A           TELEPHONE            8
  A           LOC                  30           DFT('No location assigned')
  A
```

Assume that PF1, PF2, and PF3 have the following records:

PF1

Anne	120 1st St.	USA
Doug	40 Pillsbury	Canada
Mark	2 Lakeside Dr.	Canada
Sue	120 Broadway	USA

PF2

Anne	555-1234
Doug	555-2222
Sue	555-1144

← No telephone number was found for Mark in PF2.

PF3

555-1234	Room 312
555-2222	Main lobby
999-9999	No telephone number

← No record for telephone number 555-1144 was found in PF3.

RSLH292-1

With JDFTVAL specified in the join logical file, the program reads the following logical file records:

Anne	120 1st St.	555-1234	Room 312
Doug	40 Pillsbury	555-2222	Main Lobby
Mark	2 Lakeside Dr.	999-9999	No telephone number
Sue	120 Broadway	555-1144	No location assigned

RSLH293-0

In this example, complete data is found for Anne and Doug. However, part of the data is missing for Mark and Sue.

- PF2 is missing a record for Mark because he has no telephone number. The default value for the *Telephone* field in PF2 is defined as 999-9999 using the DFT keyword. In this example, therefore, 999-9999 is the telephone number returned when no telephone number is assigned. The JDFTVAL keyword specified in the join logical file causes the default value for the *Telephone* field (which is 999-9999) in PF2 to be used to match with a record in PF3. (In PF3, a record is included to show a description for telephone number 999-9999.) Without the JDFTVAL keyword, no record would be returned for Mark.
- Sue's telephone number is not yet assigned a location; therefore, a record for 555-1144 is missing in PF3. Without JDFTVAL specified, no record would be returned for Sue. With JDFTVAL specified, the system supplies the default value specified on the DFT keyword in PF3 the *Loc* field (which is No location assigned).



## A Complex Join Logical File (Example 10)

The following example shows a more complex join logical file. Assume the data is in the following three physical files:

### Vendor Master File (PF1)

	1	2	3	4	5	6	7	8
A	R	RCD1						TEXT('VENDOR INFORMATION')
A		VDRNBR	5					TEXT('VENDOR NUMBER')
A		VDRNAM	25					TEXT('VENDOR NAME')
A		STREET	15					TEXT('STREET ADDRESS')
A		CITY	15					TEXT('CITY')
A		STATE	2					TEXT('STATE')
A		ZIPCODE	5					TEXT('ZIP CODE')
A								DFT('00000')
A		PAY	1					TEXT('PAY TERMS')
A								

### Order File (PF2)

	1	2	3	4	5	6	7	8
A	R	RCD2						TEXT('VENDORS ORDER')
A		VDRNUM	5S 0					TEXT('VENDOR NUMBER')
A		JOBNBR	6					TEXT('JOB NUMBER')
A		PRTNBR	5S 0					TEXT('PART NUMBER')
A								DFT(99999)
A		QORDER	3S 0					TEXT('QUANTITY ORDERED')
A		UNTPRC	6S 2					TEXT('PRICE')
A								

### Part File (PF3)

	1	2	3	4	5	6	7	8
A	R	RCD3						TEXT('DESCRIPTION OF PARTS')
A		PRTNBR	5S 0					TEXT('PART NUMBER')
A								DFT(99999)
A		DESCR	25					TEXT('DESCRIPTION')
A		UNITPRICE	6S 2					TEXT('UNIT PRICE')
A		WHSNBR	3					TEXT('WAREHOUSE NUMBER')
A		PRTLOC	4					TEXT('LOCATION OF PART')
A		QOHAND	5					TEXT('QUANTITY ON HAND')
A								

The join logical file record format should contain the following fields:

- Vdrnam* (vendor name)
- Street, City, State, and Zipcode* (vendor address)
- Jobnbr* (job number)
- Prtnbr* (part number)
- Descr* (description of part)
- Qorder* (quantity ordered)
- Untprc* (unit price)
- Whsnbr* (warehouse number)
- Prtloc* (location of part)

The DDS for this join logical file is as follows:



- Consider using the DYNSTL keyword. See “Dynamic Select/Omit” on page 3-12 for more details.
- Consider describing your join logical file so it can automatically share an existing access path. See “Using Existing Access Paths” on page 3-13 for more details.

**Note:** Join logical files always have access paths using the second field of the pair of fields specified in the JFLD keyword. This field acts like a key field in simple logical files. If an access path does not already exist, the access path is implicitly created with immediate maintenance.

## Data Integrity Considerations

Unless you have a lock on the physical files used by the join logical file, the following can occur:

- Your program reads a record for which there are two or more records in a secondary file. The system supplies one record to your program.
- Another program updates the record in the primary file that your program has just read, changing the join field.
- Your program issues another read request. The system supplies the next record based on the current (new) value of the join field in the primary file.

These same considerations apply to secondary files as well.

## Summary of Rules for Join Logical Files

### Requirements

The principal requirements for join logical files are:

- Each join logical file must have:
  - Only one record format, with the JFILE keyword specified for it.
  - At least two physical file names specified on the JFILE keyword. (The physical file names on the JFILE keyword do not have to be different files.)
  - At least one join specification (J in position 17 with the JFLD keyword specified).
  - A maximum of 31 secondary files.
  - At least one field name with field use other than N (neither) at the field level.
- If only two physical files are specified for the JFILE keyword, the JOIN keyword is not required. Only one join specification can be included, and it joins the two physical files.
- If more than two physical files are specified for the JFILE keyword, the following rules apply:
  - The primary file must be the first file of the pair of files specified on the first JOIN keyword (the primary file can also be the first of the pair of files specified on other JOIN keywords).

**Note:** Relative file numbers must be specified on the JOIN keyword and any JREF keyword when the same file name is specified twice on the JFILE keyword.

  - Every secondary file must be specified only once as the second file of the pair of files on the JOIN keyword. This means that for every secondary file on the JFILE keyword, one join specification must be included (two sec-

ondary files would mean two join specifications, three secondary files would mean three join specifications).

- The order in which secondary files appear in join specifications must match the order in which they are specified on the JFILE keyword.

## Join Fields

The rules to remember about join fields are:

- Every physical file you are joining must be joined to another physical file by at least one join field. A join field is a field specified as a parameter value on the JFLD keyword in a join specification.
- Join fields (specified on the JFLD keyword) must have identical attributes (length, data type, and decimal positions) or be redefined in the record format of the join logical file to have the same attributes. If the join fields are of character type, the field lengths may be different.
- Join fields need not be specified in the record format of the join logical file (unless you must redefine one or both so that their attributes are identical).
- If you redefine a join field, you can specify N in position 38 (making it a neither field) to prevent a program using the join logical file from using the redefined field.
- The maximum length of fields used in joining physical files is equal to the maximum size of keys for physical and logical files (see Appendix A, “Database File Sizes”).

## Fields in Join Logical Files

The rules to remember about fields in join logical files are:

- Fields in a record format for a join logical file must exist in one of the physical files used by the logical file or, if CONCAT, RENAME, TRNTBL, or SST is specified for the field, be a result of fields in one of the physical files.
- Fields specified as parameter values on the CONCAT keyword must be from the same physical file. If the first field name specified on the CONCAT keyword is not unique among the physical files, you must specify the JREF keyword for that field to identify which file contains the field descriptions you want to use.
- If a field name in the record format for a join logical file is specified in more than one of the physical files, you must uniquely specify on the JREF keyword which file the field comes from.
- Key fields, if specified, must come from the primary file. Key fields in the join logical file need not be key fields in the primary file.
- Select/omit fields can come from any physical file used by the join logical file, but in some circumstances the DYNSTL keyword is required.
- If specified, key fields and select/omit fields must be defined in the record format.
- Relative file numbers must be used for the JOIN and JREF keywords if the name of the physical file is specified more than once on the JFILE keyword.

## Miscellaneous

Other rules to keep in mind when using join logical files are:

- Join logical files are read-only files.
- Join record formats cannot be shared, and cannot share other record formats.
- The following are not allowed in a join logical file:
  - The REFACCPH and FORMAT keywords
  - Both fields (B specified in position 38)

---

## Chapter 4. Database Security

This chapter describes some of the database file security functions. The topics covered include database file security, public authority considerations, restricting the ability to change or delete any data in a file, and using logical files to secure data. For more information about using the security function on the AS/400 system, see *Security Reference*.

---

### File and Data Authority

The following describes the types of authority that can be granted to a user for a database file.

#### Object Operational Authority

Object operational authority is required to:

- Open the file for processing. (You must also have at least one data authority.)
- Compile a program which uses the file description.
- Display descriptive information about active members of a file.
- Open the file for query processing. For example, the Open Query File (OPNQRYF) command opens a file for query processing.

**Note:** You must also have the appropriate data authorities required by the options specified on the open operation.

#### Object Existence Authority

Object existence authority is required to:

- Delete the file.
- Save, restore, and free the storage of the file. If the object existence authority has not been explicitly granted to the user, the \*SAVSYS special user authority allows the user to save, restore, and free the storage of a file. \*SAVSYS is not the same as object existence authority.
- Remove members from the file.
- Transfer ownership of the file.

**Note:** All these functions except save/restore also require object operational authority to the file.

#### Object Management Authority

Object management authority is required to:

- Create a logical file with a keyed sequence access path (object management authority is required for the physical file referred to by the logical file).
- Grant and revoke authority. You can grant and revoke only the authority that you already have. (You must also have object operational authority to the file.)
- Change the file.
- Add members to the file. (The owner of the file becomes the owner of the new member.)
- Change the member in the file.
- Move the file.
- Rename the file.
- Rename a member of the file.

- Clear a member of the file. (Delete data authority is also required.)
- Initialize a member of the file. (Add data authority is also required to initialize with default records; delete data authority is required to initialize with deleted records.)
- Reorganize a member of the file. (All data authorities are also required.)

## Data Authorities

Data authorities can be granted only to physical files. Logical files use the data authority granted to the physical files associated with the logical file.

**Read Authority:** You can read the records in the file.

**Add Authority:** You can add new records to the file.

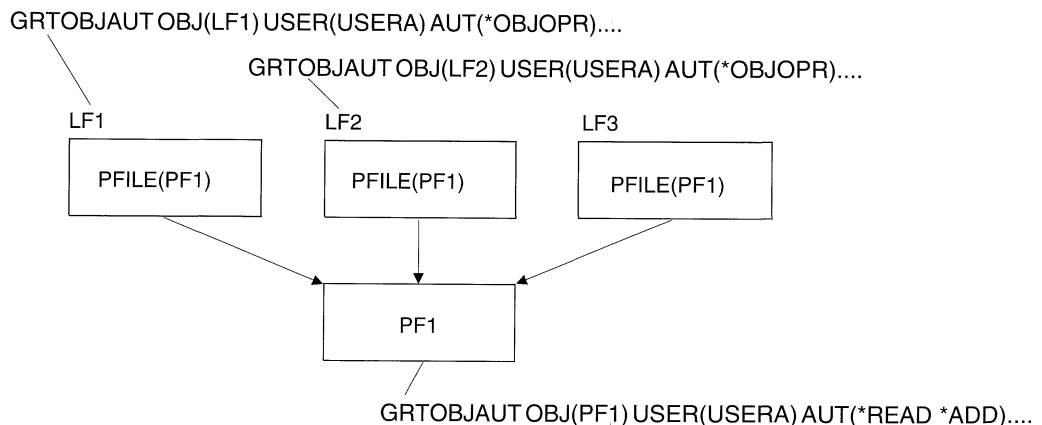
**Update Authority:** You can update existing records. (To read a record for update, you must also have read authority.)

**Delete Authority:** You can delete existing records. (To read a record for deletion, you must also have read authority.)

Normally, the authority you have to the data in the file is not verified until you actually perform the input/output operation. However, the Open Query File (OPNQRYF) and Open Database File (OPNDBF) commands also verify data authority when the file is opened.

For a physical file, authority to the file and to the data in the file can be granted. For a logical file, authority to the file can be granted, but because logical files contain no data, data authority cannot be granted. If object operational authority is not granted to a user for a physical file, that user cannot open the physical file.

The following example shows the relationship between authority granted for logical files and the physical files used by the logical file. The logical files LF1, LF2, and LF3 are based on the physical file PF1. USERA has read (\*READ) and add (\*ADD) authority to the data in PF1 and object operational authority for LF1 and LF2. This means that USERA cannot open PF1 or use its data directly in any way because the user does *not* have object operational authority (\*OBJOPR) to PF1; USERA can open LF1 and LF2 and read records from and add records to PF1 through LF1 and LF2. Note that the user was not given authority for LF3 and, therefore, cannot use it.



RV2F648-0

---

## Public Authority

When you create a file, you can specify public authority through the AUT parameter on the create command. **Public authority** is authority available to any user who does not have specific authority to the file or who is not a member of a group that has specific authority to the file. Public authority is the last authority check made. That is, if the user has specific authority to a file or the user is a member of a group with specific authority, then the public authority is not checked. Public authority can be specified as:

- \*LIBCRTAUT. The library in which the file is created is checked to determine the public authority of the file when the file is created. An authority is associated with each library. This authority is specified when the library is created, and all files created into the library are given this public authority if the \*LIBCRTAUT value is specified for the AUT parameter of the Create File (CRTLF, CRTPF, and CRTSRCPF) commands. The \*LIBCRTAUT value is the default public authority.
- \*CHANGE. All users that do not have specific user or group authority to the file have authority to change data in the file.
- \*USE. All users that do not have specific user or group authority to the file have authority to read data in the file.
- \*EXCLUDE. Only the owner, security officer, users with specific authority, or users who are members of a group with specific authority can use the file.
- \*ALL. All users that do not have specific user or group authority to the file have all data authorities along with object operational, object management, and object existence authorities.
- Authorization list name. An authorization list is a list of users and their authorities. The list allows users and their different authorities to be grouped together.

**Note:** When creating a logical file, no data authorities are granted. Consequently, \*CHANGE is the same as \*USE, and \*ALL does not grant any data authority.

You can use the Edit Object Authority (EDTOBJAUT), Grant Object Authority (GRTOBJAUT), or Revoke Object Authority (RVKOBJAUT) commands to grant or revoke the public authority of a file.

---

## Database File Capabilities

File capabilities are used to control which input/output operations are allowed for a database file independent of database file authority.

When you create a physical file, you can specify if the file is update-capable and delete-capable by using the ALWUPD and ALWDLT parameters on the Create Physical File (CRTPF) and Create Source Physical File (CRTSRCPF) commands. By creating a file that is not update-capable and not delete-capable, you can effectively enforce an environment where data cannot be changed or deleted from a file once the data is written.

File capabilities cannot be explicitly set for logical files. The file capabilities of a logical file are determined by the file capabilities of the physical files it is based on.

You cannot change file capabilities after the file is created. You must delete the file then recreate it with the desired capability. The Display File Description (DSPFD) command can be used to determine the capabilities of a file.

---

## Using Logical Files to Secure Data

You can use a logical file to prevent a field in a physical file from being viewed. This is accomplished by describing a logical file record format that does not include fields you do not want the user to see. For more information about this subject, see “Describing Logical File Record Formats” on page 3-1.

You can also use a logical file to prevent one or more fields from being changed in a physical file by specifying, for those fields you want to protect, an I (input only) in position 38 of the *DDS Coding Form*. For more information about this subject, see “Describing Field Use for Logical Files” on page 3-3.

You can use a logical file to secure records in a physical file based on the contents of one or more fields in that record. To secure records based on the contents of a field, use the select and omit keywords when describing the logical file. For more information about this subject, see “Selecting and Omitting Records Using Logical Files” on page 3-8.



---

## Part 2. Processing Database Files in Programs

The chapters in this part include information on processing database files in your programs. This includes information on planning how the file will be used in the program or job and improving the performance of your program. Descriptions of the file processing parameters and run-time options that you can specify for more efficient file processing are included in this section.

Another topic covered in this part is sharing database files across jobs so that they can be accessed by many users at the same time. Locks on files, records, or members that can prevent them from being shared across jobs are also discussed.

Using the Open Query File (OPNQRYF) command and the Open Database File (OPNDBF) command to open database file members in a program is discussed. Examples, performance considerations, and guidelines to follow when writing a high-level language program are also included. Also, typical errors that can occur are discussed.

Finally, basic database file operations are discussed. This discussion includes setting a position in the database file, and reading, updating, adding, and deleting records in a database file. A description of several ways to read database records is also included. Information on updating discusses how to change an existing database record in a logical or physical file. Information on adding a new record to a physical database member using the write operation is included. This section also includes ways you can close a database file when your program completes processing a database file member, disconnecting your program from the file. Messages to monitor when handling database file errors in a program are also discussed.



---

## Chapter 5. Run Time Considerations

Before a file is opened for processing, you should consider how you will use the file in the program and job. A better understanding of the run-time file processing parameters can help you avoid unexpected results. In addition, you might improve the performance of your program.

When a file is opened, the attributes in the database file description are merged with the parameters in the program. Normally, most of the information the system needs for your program to open and process the file is found in the file attributes and in the application program itself.

Sometimes, however, it is necessary to override the processing parameters found in the file and in the program. For example, if you want to process a member of the file other than the first member, you need a way to tell the system to use the member you want to process. The Override with Database File (OVRDBF) command allows you to do this. The OVRDBF command also allows you to specify processing parameters that can improve the performance of your job, but that cannot be specified in the file attributes or in the program. The OVRDBF command parameters take precedence over the file and program attributes. For more information on how overrides behave in the Integrated Language Environment see Appendix A of the *Integrated Language Environment\* Concepts*, SC09-1524

This chapter describes the file processing parameters. The parameter values are determined by the high-level language program, the file attributes, and any open or override commands processed before the high-level language program is called. A summary of these parameters and where you specify them can be found in "Run Time Summary" on page 5-19. For more information about processing parameters from commands, see the *CL Reference* manual for the following commands:

- Create Physical File (CRTPF)
- Create Logical File (CRTLFL)
- Create Source Physical File (CRTSRCPF)
- Add Physical File Member (ADDPFM)
- Add Logical File Member (ADDLFM)
- Change Physical File (CHGPF)
- Change Physical File Member (CHGPFM)
- Change Logical File (CHGLF)
- Change Logical File Member (CHGLFM)
- Change Source Physical File (CHGSRCPF)
- Override with Database File (OVRDBF)
- Open Database File (OPNDBF)
- Open Query File (OPNQRYF)
- Close File (CLOF)

---

### File and Member Name

**FILE and MBR Parameter.** Before you can process data in a database file, you must identify which file and member you want to use. Normally, you specify the file name and, optionally, the member name in your high-level language program. The system then uses this name when your program requests the file to be opened. To override the file name specified in your program and open a different file, you can use the TOFILE parameter on the Override with Database File (OVRDBF)

command. If no member name is specified in your program, the first member of the file (as defined by the creation date and time) is processed.

If the member name cannot be specified in the high-level language program (some high-level languages do not allow a member name), or you want a member other than the first member, you can use an Override with Database File (OVRDBF) command or an open command (OPNDBF or OPNQRYF) to specify the file and member you want to process (using the FILE and MBR parameters).

To process all the members of a file, use the OVRDBF command with the MBR(\*ALL) parameter specified. For example, if FILEX has three members and you want to process all the members, you can specify:

```
OVRDBF FILE(FILEX) MBR(*ALL)
```

If you specify MBR(\*ALL) on the OVRDBF command, your program reads the members in the order they were created. For each member, your program reads the records in keyed or arrival sequence, depending on whether the file is an arrival sequence or keyed sequence file.

---

## File Processing Options

The following section describes several run-time processing options, including identifying the file operations used by the program, specifying the starting file position, reusing deleted records, ignoring the keyed sequence access path, specifying how to handle end-of-file processing, and identifying the length of the record in the file.

### Specifying the Type of Processing

**OPTION Parameter.** When you use a file in a program, the system needs to know what types of operations you plan to use for that file. For example, the system needs to know if you plan to just read data in the file or if you plan to read and update the data. The valid operation options are: input, output, update, and delete. The system determines the options you are using from information you specify in your high-level language program or from the OPTION parameter on the Open Database File (OPNDBF) and Open Query File (OPNQRYF) commands.

The system uses the options to determine which operations are allowed in your program. For example, if you open a file for input only and your program tries an output operation, your program receives an error.

Normally, the system verifies that you have the required data authority when you do an input/output operation in your program. However, when you use the Open Query File (OPNQRYF) or Open Database File (OPNDBF) commands, the system verifies at the time the file is opened that you have the required data authority to perform the operations specified on the OPTION parameter. For more information about data authority, see “Data Authorities” on page 4-2.

The system also uses these options to determine the locks to use to protect the data integrity of the files and records being processed by your program. For more information on locks, see “Sharing Database Files Across Jobs” on page 5-6.

## Specifying the Initial File Position

**POSITION Parameter.** The system needs to know where it should start processing the file after it is opened. The default is to start just before the first record in the file (the first sequential read operation will read the first record). But, you can tell the system to start at the end of the file, or at a certain record in the middle of the file using the Override with Database File (OVRDBF) command. You can also dynamically set a position for the file in your program. For more information on setting position for a file in a program, see “Setting a Position in the File” on page 7-1.

## Reusing Deleted Records

**REUSEDLT Parameter.** When you specify REUSEDLT(\*YES) on the Create Physical File (CRTPF) or Change Physical File (CHGPF) command, the following operations may work differently:

- Arrival order becomes meaningless for a file that reuses deleted record space. Records might not be added at the end of the file.
- End-of-file delay does not work for files that reuse deleted record space.
- Applications that use DDM from a previous release system to a current release system may get different results when accessing files where deleted record space is reused.
- One hundred percent reuse of deleted record space is not guaranteed. A *file full* condition may be reached or the file may be extended even though deleted record space still exists in the file.

Because of the way the system reuses deleted record space, consider the following points before creating or changing a file to reuse deleted record space:

- Files processed using relative record numbers and files used by an application to determine a relative record number that is used as a key into another file should not reuse deleted record space.
- Files used as queues should not reuse deleted record space.
- Any files used by applications that assume new record inserts are at the end of the file should not reuse deleted record space.

If you decide to change an existing physical file to reuse deleted record space, and there are logical files with access paths with LIFO or FIFO duplicate key ordering over the physical file, you can re-create the logical files without the FIFO or LIFO attribute and avoid rebuilding the existing access path by doing the following:

1. Rename the existing logical file that has the FIFO or LIFO attribute.
2. Create a second logical file identical to the renamed file except that duplicate key ordering should not be specified for the file. Give the new file the original file name. The new file shares the access path of the renamed file.
3. Delete the renamed file.

## Ignoring the Keyed Sequence Access Path

**ACCPH Parameter.** When you process a file with a keyed sequence access path, you normally want to use that access path to retrieve the data. The system automatically uses the keyed sequence access path if a key field is defined for the file. However, sometimes you can achieve better performance by ignoring the keyed sequence access path and processing the file in arrival sequence.

You can tell the system to ignore the keyed sequence access path in some high-level languages, or on the Open Database File (OPNDBF) command. When you ignore the keyed sequence access path, operations that read data by key are not allowed. Operations are done sequentially along the arrival sequence access path. (If this option is specified for a logical file with select/omit values defined, the arrival sequence access path is used and only those records meeting the select/omit values are returned to the program. The processing is done as if the DYNSTL keyword was specified for the file.)

**Note:** You cannot ignore the keyed sequence access path for logical file members that are based on more than one physical file member.

## Delaying End of File Processing

**EOFDLY Parameter.** When you are reading a database file and your program reaches the end of the data, the system normally signals your program that there is no more data to read. Occasionally, instead of telling the program there is no more data, you might want the system to hold your program until more data arrives in the file. When more data arrives in the file, the program can read the newly arrived records. If you need that type of processing, you can use the EOFDLY parameter on the Override with Database File (OVRDBF) command. For more information on this parameter, see “Waiting for More Records When End of File Is Reached” on page 7-5.

**Note:** End of file delay should not be used for files that reuse deleted records.

## Specifying the Record Length

The system needs to know the length of the record your program will be processing, but you do not have to specify record length in your program. The system automatically determines this information from the attributes and description of the file named in your program. However, as an option, you can specify the length of the record in your high-level language program.

If the file that is opened contains records that are longer than the length specified in the program, the system allocates a storage area to match the file member's record length and this option is ignored. In this case, the entire record is passed to the program. (However, some high-level languages allow you to access only that portion of the record defined by the record length specified in the program.) If the file that is opened contains records that are less than the length specified in the program, the system allocates a storage area for the program-specified record length. The program can use the extra storage space, but only the record lengths defined for the file member are used for input/output operations.

## Ignoring Record Formats

When you use a multiple format logical file, the system assumes you want to use all formats defined for that file. However, if you do not want to use all of the formats, you can specify which formats you want to use and which ones you want to ignore. If you do not use this option to ignore formats, your program can process all formats defined in the file. For more information about this processing option, see your high-level language guide.

## Determining If Duplicate Keys Exist

**DUPKEYCHK Parameter.** The set of keyed sequence access paths used to determine if the key is a duplicate key differs depending on the I/O operation that is performed.

For input operations (reads), the keyed sequence access path used is the one that the file is opened with. Any other keyed sequence access paths that can exist over the physical file are not considered. Also, any records in the keyed sequence access path omitted because of select/omit specifications are not considered when deciding if the key operation is a duplicate.

For output (write) and update operations, all nonunique keyed sequence access paths of \*IMMED maintenance that exist over the physical file are searched to determine if the key for this output or update operation is a duplicate. Only keyed sequence access paths that have \*RBLD and \*DLY maintenance are considered if the access paths are actively open over the file at feedback time.

When you process a keyed file with a COBOL program, you can specify duplicate key feedback to be returned to your program through the COBOL language, or on the Open Database File (OPNDBF) or Open Query File (OPNQRYF) commands. However, in COBOL having duplicate key feedback returned can cause a decline in performance.

---

## Data Recovery and Integrity

The following section describes data integrity run-time considerations.

## Protecting Your File with the Journaling and Commitment Control

**COMMIT Parameter.** Journaling and commitment control are the preferred methods for data and transaction recovery on the AS/400 system. Database file journaling is started by running the Start Journaling Physical File (STRJRNPF) command for the file. Access path journaling is started by running the Start Journaling Access Path (STRJRNAP) command for the file. You tell the system that you want your files to run under commitment control through the Start Commitment Control (STRCMTCTL) command and through high-level language specifications. You can also specify the commitment control (COMMIT) parameter on the Open Database File (OPNDBF) and Open Query File (OPNQRYF) commands. For more information on journaling and commitment control, see Chapter 13, "Database Recovery Considerations," and the *Advanced Backup and Recovery Guide*.

## Writing Data and Access Paths to Auxiliary Storage

**FRCRATIO and FRCACPTH Parameters.** Normally, the AS/400 integrated database management system determines when to write changed data from main storage to auxiliary storage. If you want to control when database changes are written to auxiliary storage, you can use the force write ratio (FRCRATIO) parameter on either the create, change, or override database file commands, and the force access path (FRCACPTH) parameter on the create and change database file commands. Using the FRCRATIO and FRCACPTH parameters have performance and recovery considerations for your system. To understand these considerations, see Chapter 13, “Database Recovery Considerations.”

## Checking Changes to the Record Format Description

**LVLCHK Parameter.** The system checks, when you open the file, if the description of the record format you are using was changed since the program was compiled to an extent that your program cannot process the file. The system normally notifies your program of this condition. This condition is known as a level check. When you use the create or change file commands, you can specify that you want level checking. You can also override the level check attribute defined for the file using the LVLCHK parameter on the Override with Database File (OVRDBF) command. For more information about this parameter, see “Effect of Changing Fields in a File Description” on page 11-1.

## Checking for the File’s Expiration Date

**EXPDATE and EXPCHK Parameters.** The system can verify that the data in the file you specify is still current. You can specify the expiration date for a file or member using the EXPDATE parameter on the create and change file commands, and you can specify whether or not the system is to check that date using the EXPCHK parameter on the Override with Database File (OVRDBF) command. If you do check the expiration date and the current date is greater than the expiration date, a message is sent to the system operator when the file is opened.

## Preventing the Job from Changing Data in the File

**INHWRT Parameter.** If you want to test your program, but do not want to actually change data in files used for the test, you can tell the system to not write (inhibit) any changes to the file that the program attempts to make. To inhibit any changes to the file, specify INHWRT(\*YES) on the Override with Database File (OVRDBF) command.

---

## Sharing Database Files Across Jobs

By definition, all database files can be used by many users at the same time. However, some operations can lock the file, member, or data records in a member to prevent the file, member, or data records from being shared across jobs.

Each database command or high-level language program allocates the file, member, and data records that it uses. Depending on the operation requested, other users will not be able to use the allocated file, member, or records. An operation on a logical file or logical file member can allocate the file(s) and member(s) that the logical file depends on for data or an access path.



For example, the open of a logical file allocates the data of the physical file member that the logical file is based on. If the program updates the logical file member, another user may not request, at the same time, that the physical file member used by that logical file member be cleared of data.

For a list of commonly used database functions and the types of locks they place on database files, see Appendix C, "Database Lock Considerations."

## Record Locks

**WAITRCD Parameter.** The AS/400 database has built-in integrity for records. For example, if PGMA reads a record for update, it locks that record. Another program may not read the same record for update until PGMA releases the record, but another program could read the record just for inquiry. In this way, the system ensures the integrity of the database.

The system determines the lock condition based on the type of file processing specified in your program and the operation requested. For example, if your open options include update or delete, each record read is locked so that any number of users can read the record at the same time, but only one user can update the record.

The system normally waits a specific number of seconds for a locked record to be released before it sends your program a message that it cannot get the record you are requesting. The default record wait time is 60 seconds; however, you can set your own wait time through the WAITRCD parameter on the create and change file commands and the override database file command. If your program is notified that the record it wants is locked by another operation, you can have your program take the appropriate action (for example, you could send a message to the operator that the requested record is currently unavailable).

The system automatically releases a lock when the locked record is updated or deleted. However, you can release record locks without updating the record. For information on how to release a record lock, see your high-level language guide.

**Note:** Using commitment control changes the record locking rules. See the *Advanced Backup and Recovery Guide* for more information on commitment control and its effect on the record locking rules.

You can use the Display Record Locks (DSPRCDLCK) command to display the current lock status (wait or held) of records for a physical file member. The command will also indicate what type of lock is currently held. (For more information about lock types, see the *Advanced Backup and Recovery Guide*.) Depending on the parameters you specify, this command displays the lock status for a specific record or displays the lock status of all records in the member. You can also display record locks from the Work with Job (WRKJOB) display.

You can determine if your job currently has any records locked using the Check Record Lock (CHKRCDLCK) command. This command returns a message (which you can monitor) if your job has any locked records. The command is useful if you are using group jobs. For example, you could check to see if you had any records locked, before transferring to another group job. If you determined you did have records locked, your program could release those locks.

## File Locks

**WAITFILE Parameter.** Some file operations exclusively allocate the file for the length of the operation. During the time the file is allocated exclusively, any program trying to open that file has to wait until the file is released. You can control the amount of time a program waits for the file to become available by specifying a wait time on the WAITFILE parameter of the create and change file commands and the override database file command. If you do not specifically request a wait time, the system defaults the file wait time to zero seconds.

A file is exclusively allocated when an operation that changes its attributes is run. These operations (such as move, rename, grant or revoke authority, change owner, or delete) cannot be run at the same time with any other operation on the same file or on members of that file. Other file operations (such as display, open, dump, or check object) only use the file definition, and thus lock the file less exclusively. They can run at the same time with each other and with input/output operations on a member.

## Member Locks

Member operations (such as add and remove) automatically allocate the file exclusively enough to prevent other file operations from occurring at the same time. Input/output operations on the same member cannot be run, but input/output operations on other members of the same file can run at the same time.

## Record Format Data Locks

**RCDFMTLCK Parameter.** If you want to lock the entire set of records associated with a record format (for example, all the records in a physical file), you can use the RCDFMTLCK parameter on the OVRDBF command.

---

## Sharing Database Files in the Same Job or Activation Group

**SHARE Parameter.** By default, the database management system lets one file be read and changed by many users at the same time. You can also share a file in the same job or activation group by opening the database file:

- More than once in the same program.
- In different programs in the same job or activation group.

**Note:** For more information on open sharing in the Integrated Language Environment see Appendix A in the *Integrated Language Environment Concepts* manual, SC09-1524.

The SHARE parameter on the create file, change file, and override database file commands allow sharing in a job or activation group, including sharing the file, its status, its positions, and its storage area. Sharing files in the job or activation group can improve performance by reducing the amount of main storage needed and by reducing the time needed to open and close the file.

Using the SHARE(\*YES) parameter lets two or more programs running in the same job or activation group share an open data path (ODP). An open data path is the path through which all input/output operations for the file are performed. In a sense, it connects the program to a file. If you do not specify the SHARE(\*YES) parameter, a new open data path is created every time a file is opened. If an active file is opened more than once in the same job or activation group, you can

use the active ODP for the file with the current open of the file. You do not have to create a new open data path.

This reduces the amount of time required to open the file after the first open, and the amount of main storage required by the job or activation group. SHARE(\*YES) must be specified for the first open and other opens of the same file for the open data path to be shared. A well-designed (for performance) application normally shares an open data path with files that are opened in multiple programs in the same job or activation group.

Specifying SHARE(\*NO) tells the system not to share the open data path for a file. Normally, this is specified only for those files that are seldom used or require unique processing in specific programs.

**Note:** A high-level language program processes an open or a close operation as though the file were not being shared. You do not specify that the file is being shared in the high-level language program. You indicate that the file is being shared in the same job or activation group through the SHARE parameter. The SHARE parameter is specified only on the create, change, and override database file commands.

## Open Considerations for Files Shared in a Job or Activation Group

Consider the following items when you open a database file that is shared in the same job or activation group.

- Make sure that when the shared file is opened for the first time in a job or activation group, all the open options needed for subsequent opens of the file are specified. If the open options specified for subsequent opens of a shared file do not match those specified for the first open of a shared file, an error message is sent to the program. (You can correct this by making changes to your program or to the OPNDBF or OPNQUERYF command parameters, to remove any incompatible options.)

For example, PGMA is the first program to open FILE1 in the job or activation group and PGMA only needs to read the file. However, PGMA calls PGMB, which will delete records from the same shared file. Because PGMB will delete records from the shared file, PGMA will have to open the file as if it, PGMA, is also going to delete records. You can accomplish this by using the correct specifications in the high-level language. (To accomplish this in some high-level languages, you may have to use file operation statements that are never run. See your high-level language guide for more details.) You can also specify the file processing option on the OPTION parameter on the Open Database File (OPNDBF) and Open Query File (OPNQUERYF) commands.

- Sometimes sharing a file within a job or activation group is not desirable. For example, one program can need records from a file in arrival sequence and another program needs the records in keyed sequence. In this situation, you should not share the open data path. You would specify SHARE(\*NO) on the Override with Database File (OVRDBF) command to ensure the file was not shared within the job or activation group.
- If debug mode is entered with UPDPROD(\*NO) after the first open of a shared file in a production library, subsequent shared opens of the file share the original open data path and allow the file to be changed. To prevent this, specify SHARE(\*NO) on the OVRDBF command before opening files being debugged.

- The use of commitment control for the first open of a shared file requires that all subsequent shared opens also use commitment control.
- Key feedback, insert key feedback, or duplicate key feedback must be specified on the full open if any of these feedback types are desired on the subsequent shared opens of the file.
- If you did not specify a library name in the program or on the Override with Database File (OVRDBF) command (\*LIBL is used), the system assumes the library list has not changed since the last open of the same shared file with \*LIBL specified. If the library list has changed, you should specify the library name on the OVRDBF command to ensure the correct file is opened.
- The record length that is specified on the full open is the record length that is used on subsequent shared opens even if a larger record length value is specified on the shared opens of the file.
- Overrides and program specifications specified on the first open of the shared file are processed. Overrides and program specifications specified on subsequent opens, other than those that change the file name or the value specified on the SHARE or LVLCHK parameters on the OVRDBF command, are ignored.
- Overrides specified for a first open using the OPNQRYF command can be used to change the names of the files, libraries, and members that should be processed by the Open Query File command. Any parameter values specified on the Override with Database File (OVRDBF) command other than TOFILE, MBR, LVLCHK, and SEQONLY are ignored by the OPNQRYF command.
- The Open Database File (OPNDBF) and Open Query File (OPNQRYF) commands scope the ODP to the level specified on the Open Scope (OPNSCOPE) parameter according to the following:
  - The system searches for shared opens in the activation group first, and then in the job.
  - Shared opens that are scoped to an activation group may not be shared between activation groups.
  - Shared opens that are scoped to the job can be shared throughout the job, by any number of activation groups at a time.

The CPF4123 diagnostic message lists the mismatches that can be encountered between the full open and the subsequent shared opens. These mismatches do not cause the shared open to fail.

**Note:** The Open Query File (OPNQRYF) command never shares an existing shared open data path in the job or activation group. If a shared ODP already exists in the job or activation group with the same file, library, and member name as the one specified on the Open Query File command, the system sends an error message and the query file is not opened.

## Input/Output Considerations for Files Shared in a Job or Activation Group

Consider the following items when processing a database file that is shared in the same job or activation group.

- Because only one open data path is allowed for a shared file, only one record position is maintained for all programs in the job or activation group that is sharing the file. If a program establishes a position for a record using a read or a read-for-update operation, then calls another program that also uses the

shared file, the record position may have moved or a record lock been released when the called program returns to the calling program. This can cause errors in the calling program because of an unexpected record position or lock condition. When sharing files, it is your responsibility to manage the record position and record locking considerations by re-establishing position and locks.

- If a shared file is first opened for update, this does not necessarily cause every subsequent program that shares the file to request a record lock. The system determines the type of record lock needed for each program using the file. The system tries to keep lock contention to a minimum, while still ensuring data integrity.

For example, PGMA is the first program in the job or activation group to open a shared file. PGMA intends to update records in the file; therefore, when the program reads a record for update, it will lock the record. PGMA then calls PGMB. PGMB also uses the shared file, but it does not update any records in the file; PGMB just reads records. Even though PGMA originally opened the shared file as update-capable, PGMB will not lock the records it reads, because of the processing specifications in PGMB. Thus, the system ensures data integrity, while minimizing record lock contention.

## Close Considerations for Files Shared in a Job or Activation Group

Consider the following items when closing a database file that is shared in the same job or activation group.

- The complete processing of a close operation (including releasing file, member, and record locks; forcing changes to auxiliary storage; and destroying the open data path) is done only when the last program to open the shared open data path closes it.
- If the file was opened with the Open Database File (OPNDBF) or the Open Query File (OPNQRYF) command, use the Close File (CLOF) command to close the file. The Reclaim Resources (RCLRSC) command can be used to close a file opened by the Open Query File (OPNQRYF) command when one of the following is specified:
  - OPNSCOPE(\*ACTGRPDFN), and the open is requested from the default activation group.
  - TYPE(\*NORMAL) is specified.

If one of the following is specified, the file remains open even if the Reclaim Resources (RCLRSC) command is run:

- OPNSCOPE(\*ACTGRPDFN), and the open is requested from some activation group other than the default
- OPNSCOPE(\*ACTGRP)
- OPNSCOPE(\*JOB)
- TYPE(\*PERM)

### Examples of Closing Shared Files

The following examples show some of the things to consider when closing a file that is shared in the same job.

**Example 1:** Using a single set of files with similar processing options.

In this example, the user signs on and most of the programs used process the same set of files.

A CL program (PGMA) is used as the first program (to set up the application, including overrides and opening the shared files). PGMA then transfers control to PGMB, which displays the application menu. Assume, in this example, that files A, B, and C are used, and files A and B are to be shared. Files A and B were created with SHARE(\*NO); therefore an OVRDBF command should precede each of the OPNDBF commands to specify the SHARE(\*YES) option. File C was created with SHARE(\*NO) and File C is not to be shared in this example.

```
PGMA:  PGM      /* PGMA - Initial program */
       OVRDBF  FILE(A) SHARE(*YES)
       OVRDBF  FILE(B) SHARE(*YES)
       OPNDBF  FILE(A) OPTION(*ALL) ....
       OPNDBF  FILE(B) OPTION(*INP) ...
       TFRCTL  PGMB
       ENDPGM
```

```
PGMB:  PGM      /* PGMB - Menu program */
       DCLF    FILE(DISPLAY)
BEGIN:  SNDRCVF  RCDfmt(MENU)
       IF      (&RESPONSE *EQ '1') CALL PGM11
       IF      (&RESPONSE *EQ '2') CALL PGM12
       .
       .
       IF      (&RESPONSE *EQ '90') SIGNOFF
       GOTO    BEGIN
       ENDPGM
```

The files opened in PGMA are either scoped to the job, or PGMA, PGM11, and PGM12 run in the same activation group and the file opens are scoped to that activation group.

In this example, assume that:

- PGM11 opens files A and B. Because these files were opened as shared by the OPNDBF commands in PGMA, the open time is reduced. The close time is also reduced when the shared open data path is closed. The Override with Database File (OVRDBF) commands remain in effect even though control is transferred (with the Transfer Control [TFRCTL] command in PGMA) to PGMB.
- PGM12 opens files A, B, and C. File A and B are already opened as shared and the open time is reduced. Because file C is used only in this program, the file is not opened as shared.

In this example, the Close File (CLOF) was not used because only one set of files is required. When the operator signs off, the files are automatically closed. It is assumed that PGMA (the initial program) is called only at the start of the job. For information on how to reclaim resources in the Integrated Language Environment, see Appendix A in the *Integrated Language Environment Concepts* manual, SC09-1524.

**Note:** The display file (DISPLAY) in PGMB can also be specified as a shared file, which would improve the performance for opening the display file in any programs that use it later.

In Example 1, the OPNDBF commands are placed in a separate program (PGMA) so the other processing programs in the job run as efficiently as possible. That is,

the important files used by the other programs in the job are opened in PGMA. After the files are opened by PGMA, the main processing programs (PGMB, PGM11, and PGM12) can share the files; therefore, their open and close requests will process faster. In addition, by placing the open commands (OPNDBF) in PGMA rather than in PGMB, the amount of main storage used for PGMB is reduced.

Any overrides and opens can be specified in the initial program (PGMA); then, that program can be removed from the job (for example, by transferring out of it). However, the open data paths that the program created when it opened the files remain in existence and can be used by other programs in the job.

Note the handling of the OVRDBF commands in relation to the OPNDBF commands. Overrides must be specified before the file is opened. Some of the parameters on the OVRDBF command also exist on the OPNDBF command. If conflicts arise, the OVRDBF value is used. For more information on when overrides take effect in the Integrated Language Environment, see Appendix A in *Integrated Language Environment Concepts*, SC09-1524.

**Example 2:** Using multiple sets of files with similar processing options.

Assume that a menu requests the operator to specify the application program (for example, accounts receivable or accounts payable) that uses the Open Database File (OPNDBF) command to open the required files. When the application is ended, the Close File (CLOF) command closes the files. The CLOF command is used to help reduce the amount of main storage needed by the job. In this example, different files are used for each application. The user normally works with one application for a considerable length of time before selecting a new application.

An example of the accounts receivable programs follows:

```

PGMC:  PGM      /* PGMC PROGRAM */
       DCLF     FILE(DISPLAY)
BEGIN: SNDRCVF  RCDFMT(TOPMENU)
       IF       (&RESPONSE *EQ '1') CALL ACCRECV
       IF       (&RESPONSE *EQ '2') CALL ACCPAY
       .
       .
       IF       (&RESPONSE *EQ '90') SIGNOFF
       GOTO     BEGIN
       ENDPGM

ACCREC: PGM      /* ACCREC PROGRAM */
       DCLF     FILE(DISPLAY)
       OVRDBF   FILE(A) SHARE(*YES)
       OVRDBF   FILE(B) SHARE(*YES)
       OPNDBF   FILE(A) OPTION(*ALL) ....
       OPNDBF   FILE(B) OPTIONS(*INP) ...
BEGIN: SNDRCVF  RCDFMT(ACCRMENU)
       IF       (&RESPONSE *EQ '1') CALL PGM21
       IF       (&RESPONSE *EQ '2') CALL PGM22
       .
       .
       IF       (&RESPONSE *EQ '88') DO /* Return */
           CLOF FILE(A)
           CLOF FILE(B)
           RETURN
           ENDDO
       GOTO     BEGIN
       ENDPGM

```

The program for the accounts payable menu would be similar, but with a different set of OPNDBF and CLOF commands.

For this example, files A and B were created with SHARE(\*NO). Therefore, an OVRDBF command must precede the OPNDBF command. As in Example 1, the amount of main storage used by each job could be reduced by placing the OPNDBF commands in a separate program and calling it. A separate program could also be created for the CLOF commands. The OPNDBF commands could be placed in an application setup program that is called from the menu, which transfers control to the specific application program menu (any overrides specified in this setup program are kept). However, calling separate programs for these functions also uses system resources and, depending on the frequency with which the different menus are used, it can be better to include the OPNDBF and CLOF commands in each application program menu as shown in this example.

Another choice is to use the Reclaim Resources (RCLRSC) command in PGMC (the setup program) instead of using the Close File (CLOF) commands. The RCLRSC command closes any files and frees any leftover storage associated with any files and programs that were called and have since returned to the calling program. However, RCLRSC does *not* close files that are opened with the following specified on the Open Database File (OPNDBF) or Open Query File (OPNQRYF) commands:

- OPNSCOPE(\*ACTGRPDFN), and the open is requested from some activation group other than the default.



|  
|  
|  
|  
|

- OPNSCOPE(\*ACTGRP) reclaims if the RCLRSC command is from an activation group with an activation group number that is lower than the activation group number of the open.
- OPNSCOPE(\*JOB).
- TYPE(\*PERM).

The following example shows the RCLRSC command used to close files:

```
.  
.  
IF          (&RESPONSE *EQ '1') DO  
            CALL ACCRECV  
            RCLRSC  
            ENDDO  
IF          (&RESPONSE *EQ '2') DO  
            CALL ACCPAY  
            RCLRSC  
            ENDDO  
.  
.
```

**Example 3:** Using a single set of files with different processing requirements.

If some programs need read-only file processing and others need some or all of the options (input/update/add/delete), one of the following methods can be used. The same methods apply if a file is to be processed with certain command parameters in some programs and not in others (for example, sometimes the commit option should be used).

A single Open Database File (OPNDBF) command could be used to specify OPTION(\*ALL) and the open data path would be opened shared (if, for example, a previous OVRDBF command was used to specify SHARE(\*YES)). Each program could then open a subset of the options. The program requests the type of open depending on the specifications in the program. In some cases this does not require any more considerations because a program specifying an open for input only would operate similarly as if it had not done a shared open (for example, no additional record locking occurs when a record is read).

However, some options specified on the OPNDBF command can affect how the program operates. For example, SEQONLY(\*NO) is specified on the open command for a file in the program. An error would occur if the OPNDBF command used SEQONLY(\*YES) and a program attempted an operation that was not valid with sequential-only processing.

The ACCPTH parameter must also be consistent with the way programs will use the access path (arrival or keyed).

If COMMIT(\*YES) is specified on the Open Database File (OPNDBF) command and the Start Commitment Control (STRCMTCTL) command specifies LCKLVL(\*ALL) or LCKLVL(\*CS), any read operation of a record locks that record (per commitment control record locking rules). This can cause records to be locked unexpectedly and cause errors in the program.

Two OPNDBF commands could be used for the same data (for example, one with OPTION(\*ALL) and the other specifying OPTION(\*INP)). The second use must be

a logical file pointing to the same physical file(s). This logical file can then be opened as SHARE(\*YES) and multiple uses made of it during the same job.

---

## Sequential-Only Processing

**SEQONLY and NBRRCDs Parameters.** If your program processes a database file sequentially for input only or output only, you might be able to improve performance using the sequential-only processing (SEQONLY) parameter on the Override with Database File (OVRDBF) or the Open Database File (OPNDBF) commands. To use SEQONLY processing, the file must be opened for input-only or output-only. The NBRRCDs parameter can be used with any combination of open options. (The Open Query File [OPNQRYF] command uses sequential-only processing whenever possible.) Depending on your high-level language specifications, the high-level language can also use sequential-only processing as the default. For example, if you open a file for input only and the only file operations specified in the high-level language program are sequential read operations, then the high-level language automatically requests sequential-only processing.

**Note:** File positioning operations are not considered sequential read operations; therefore, a high-level language program containing positioning operations will *not* automatically request sequential-only processing. (The SETLL operation in the RPG/400 language and the START operation in the COBOL/400\* language are examples of file positioning operations.) Even though the high-level language program can not automatically request sequential-only processing, you can request it using the SEQONLY parameter on the OVRDBF command.

If you specify sequential-only processing, you can also specify the number of records to be moved as one unit between the system database main storage area and the job's internal data main storage area. If you do not specify the sequential-only number of records to be moved, the system calculates a number based on the number of records that fit into a 4096-byte buffer.

The system also provides you a way to control the number of records that are moved as a unit between auxiliary storage and main storage. If you are reading the data in the file in the same order as the data is physically stored, you can improve the performance of your job using the NBRRCDs parameter on the OVRDBF command.

**Note:** Sequential-only processing should not be used with a keyed sequence access path file unless the physical data is in the same order as the access path. SEQONLY(\*YES) processing may cause poor application performance until the physical data is reorganized into the access path's order.

## Open Considerations for Sequential-Only Processing

The following considerations apply for opening files when sequential-only processing is specified. If the system determines that sequential-only processing is not allowed, a message is sent to the program to indicate that the request for sequential-only processing is not being accepted; however, the file is still opened for processing.

- If the program opened the member for output only, and if SEQONLY(\*YES) was specified (number of records was not specified) and either the opened member is a logical member, a uniquely keyed physical member, or there are other access paths to the physical member, SEQONLY(\*YES) is changed to

SEQONLY(\*NO) so the program can handle possible errors (for example, duplicate keys, conversion mapping, and select/omit errors) at the time of the output operation. If you want the system to run sequential-only processing, change the SEQONLY parameter to include both the \*YES value and number of records specification.

- Sequential-only processing can be specified only for input-only (read) or output-only (add) operations. If the program specifies update or delete operations, sequential-only processing is not allowed by the system.
- If a file is being opened for output, it must be a physical file or a logical file based on one physical file member.
- Sequential-only processing can be specified with commitment control only if the member is opened for output-only.
- If sequential-only processing is being used for files opened with commitment control and a rollback operation is performed for the job, the records that reside in the job's storage area at the time of the rollback operation are not written to the system storage area and never appear in the journal for the commitment control transaction. If no records were ever written to the system storage area prior to a rollback operation being performed for a particular commitment control transaction, the entire commitment control transaction is not reflected in the journal.
- For output-only, the number of records specified to be moved as a unit and the force ratio are compared and automatically adjusted as necessary. If the number of records is larger than the force ratio, the number of records is reduced to equal the force ratio. If the opposite is true, the force ratio is reduced to equal the number of records.
- If the program opened the member for output only, and if SEQONLY(\*YES) was specified (number of records was not specified), and duplicate or insert key feedback has been requested, SEQONLY(\*YES) will be changed to SEQONLY(\*NO) to provide the feedback on a record-by-record basis when the records are inserted into the file.
- The number of records in a block will be changed to one if all of the following are true:
  - The member was opened for output-only processing.
  - No override operations are in effect that have specified sequential-only processing.
  - The file being opened is a file that cannot be extended because its increment number of records was set to zero.
  - The number of bytes available in the file is less than the number of bytes that fit into a block of records.

The following considerations apply when sequential-only processing is not specified and the file is opened using the Open Query File (OPNQRYF) command. If these conditions are satisfied, a message is sent to indicate that sequential-only processing will be performed and the query file is opened.

- If the OPNQRYF command specifies the name of one or more fields on the group field (GRPFLD) parameter, or OPNQRYF requires group processing.
- If the OPNQRYF command specifies one or more fields, or \*ALL on the UNIQUEKEY parameter.

- If a view is used with the DISTINCT option on the SQL/400 SELECT statement, then SEQONLY(\*YES) processing is automatically performed.

For more details about the OPNQRYF command, see “Using the Open Query File (OPNQRYF) Command” on page 6-3.

## Input/Output Considerations for Sequential-Only Processing

The following considerations apply for input/output operations on files when sequential-only processing is specified.

- For input, your program receives one record at a time from the input buffer. When all records in the input buffer are processed, the system automatically reads the next set of records.  
**Note:** Changes made after records are read into the input buffer are not reflected in the input buffer.
- For output, your program must move one record at a time to the output buffer. When the output buffer is full, the system automatically adds the records to the database.  
**Note:** If you are using a journal, the entire buffer is written to the journal at one time as if the entries had logically occurred together.

If you use sequential-only processing for output, you might not see all the changes made to the file as they occur. For example, if sequential-only is specified for a file being used by PGMA, and PGMA is adding new records to the file and the SEQONLY parameter was specified with 5 as the number of records in the buffer, then only when the buffer is filled will the newly added records be transferred to the database. In this example, only when the fifth record was added, would the first five records be transferred to the database, and be available for processing by other jobs in the system.

In addition, if you use sequential-only processing for output, some additions might not be made to the database if you do not handle the errors that could occur when records are moved from the buffer to the database. For example, assume the buffer holds five records, and the third record in the buffer had a key that was a duplicate of another record in the file and the file was defined as a unique-key file. In this case, when the system transfers the buffer to the database it would add the first two records and then get a duplicate key error for the third. Because of this error, the third, fourth, and fifth records in the buffer would *not* be added to the database.

- The force-end-of-data function can be used for output operations to force all records in the buffer to the database (except those records that would cause a duplicate key in a file defined as having unique keys, as described previously). The force-end-of-data function is only available in certain high-level languages.
- The number of records in a block will be changed to one if all of the following are true:
  - The member was opened for output-only processing or sequential-only processing.
  - No override operations are in effect that have specified sequential-only processing.
  - The file being opened is being extended because the increment number of records was set to zero.

- The number of bytes available in the file is less than the number of bytes that fit into a block of records.

## Close Considerations for Sequential-Only Processing

When a file for which sequential-only processing is specified is closed, all records still in the output buffer are added to the database. However, if an error occurs for a record, any records following that record are not added to the database.

If multiple programs in the same job are sharing a sequential-only output file, the output buffer is not emptied until the final close occurs. Consequently, a close (other than the last close in the job) does not cause the records still in the buffer to appear in the database for this or any other job.

---

## Run Time Summary

The following tables list parameters that control your program's use of the database file member, and indicates where these parameters can be specified. For parameters that can be specified in more than one place, the system merges the values. The Override with Database File (OVRDBF) command parameters take precedence over program parameters, and Open Database File (OPNDBF) or Open Query File (OPNQRYF) command parameters take precedence over create or change file parameters.

**Note:** Any override parameters other than TOFILE, MBR, LVLCHK, SEQONLY, SHARE, WAITRCD, and INHWRT are ignored by the OPNQRYF command.

A table of database processing options specified on control language (CL) commands is shown below:

Figure 5-1 (Page 1 of 2). Database Processing Options Specified on CL Commands

Description	Parameter	CRTPF or CRTLF Cmd	CHGPF or CHGLF Cmd	OPNDBF Cmd	OPNQRYF Cmd	OVRDBF Cmd
File name	FILE	X	X <sup>1</sup>	X	X	X
Library name		X	X <sup>2</sup>	X	X	X
Member name	MBR	X		X	X	X
Member processing options	OPTION			X	X	
Record format lock state	RCDFMTLCK					X
Starting file position after open	POSITION					X
Program performs only sequential processing	SEQONLY			X	X	X
Ignore keyed sequence access path	ACCPTH			X		
Time to wait for file locks	WAITFILE	X	X			X

Figure 5-1 (Page 2 of 2). Database Processing Options Specified on CL Commands

Description	Parameter	CRTPF or CRTLF Cmd	CHGPF or CHGLF Cmd	OPNDBF Cmd	OPNQRYF Cmd	OVRDBF Cmd
Time to wait for record locks	WAITRCD	X	X			X
Prevent overrides	SECURE					X
Number of records to be transferred from auxiliary to main storage	NBRRCD					X
Share open data path with other programs	SHARE	X	X			X
Format selector	FMTSLR	X <sup>3</sup>	X <sup>3</sup>			X
Force ratio	FRCRATIO	X	X			X
Inhibit write	INHWRT					X
Level check record formats	LVLCHK	X	X			X
Expiration date checking	EXPCHK					X
Expiration date	EXPDATE	X <sup>4</sup>	X <sup>4</sup>			X
Force access path	FRCACCPH	X	X			
Commitment control	COMMIT			X	X	
End-of-file delay	EOFDLY					X
Duplicate key check	DUPKEYCHK			X	X	
Reuse deleted record space	REUSEDLT	X <sup>4</sup>	X <sup>4</sup>			
Coded character set identifier	CCSID	X <sup>4</sup>	X <sup>4</sup>			
Sort Sequence	SRTSEQ	X	X		X	
Language identifier	LANGID	X	X		X	

- 1 File name: The CHGPF and CHGLF commands use the file name for identification only. You cannot change the file name.
- 2 Library name: The CHGPF and CHGLF commands use the library name for identification only. You cannot change the library name.
- 3 Format selector: Used on the CRTLF and CHGLF commands only.
- 4 Expiration date, reuse deleted records, and coded character set identifier: Used on the CRTPF and CHGPF commands only.

A table of database processing options specified in programs is shown below:

Figure 5-2. Database Processing Options Specified in Programs

Description	RPG/400 Language	COBOL/400 Language	AS/400 BASIC	AS/400 PL/I	AS/400 Pascal
File name	X	X	X	X	X
Library name			X	X	X
Member name			X	X	X
Program record length	X	X	X	X	X
Member processing options	X	X	X	X	X
Record format lock state			X	X	
Record formats the program will use	X		X		
Clear physical file member of records		X	X		X
Program performs only sequential processing	X	X		X	X
Ignore keyed sequence access path	X	X	X	X	X
Share open data path with other programs				X	X
Level check record formats	X	X	X	X	
Commitment control	X	X		X	
Duplicate key check		X			

**Note:** Control language (CL) programs can also specify many of these parameters. See Figure 5-1 on page 5-19 for more information about the database processing options that can be specified on CL commands.

## Storage Pool Paging Option Effect on Database Performance

The Paging option of shared pools can have a significant impact on the performance of reading and changing database files.

- A paging option of \*FIXED causes the program to minimize the amount of memory it uses by:
  - Transferring data from auxiliary storage to main memory in smaller blocks
  - Writing file changes (updates to existing records or newly added records) to auxiliary storage frequently

This option allows the system to perform much like it did before the paging option was added.

- A paging option of \*CALC may improve how the program performs when it reads and updates database files. In cases where there is sufficient memory available within a shared pool, the program may:
  - Transfer larger blocks of data to memory from auxiliary storage.
  - Write changed data to auxiliary storage less frequently.

The paging operations done on database files vary dynamically based on file use and memory availability. Frequently referenced files are more likely to remain resident than those less often accessed. The memory is used somewhat like a cache for popular data. The overall number of I/O operations may be reduced using the \*CALC paging option.

For more information on the paging option see the “Automatic System Tuning” section of the *Work Management Guide*.





---

## Chapter 6. Opening a Database File

This chapter discusses opening a database file. In addition, the CL commands Open Database File (OPNDBF) and Open Query File (OPNQRYF) are discussed.

---

### Opening a Database File Member

To use a database file in a program, your program must issue an open operation to the database file. If you do not specify an open operation in some programming languages, they automatically open the file for you. If you did not specify a member name in your program or on an Override with Database File (OVRDBF) command, the first member (as defined by creation date and time) in the file is used.

If you specify a member name, files that have the correct file name but do not contain the member name are ignored. If you have multiple database files named FILEA in different libraries, the member that is opened is the first one in the library list that matches the request. For example, LIB1, LIB2, and LIB3 are in your library list and all three contain a file named FILEA. Only FILEA in LIB3 has a member named MBRA that is to be opened. Member MRBA in FILEA in LIB3 is opened; the other FILEAs are ignored.

After finding the member, the system connects your program to the database file. This allows your program to perform input/output operations to the file. For more information about opening files in your high-level language program, see the appropriate high-level language guide.

You can open a database file with statements in your high-level language program. You can also use the CL open commands: Open Database File (OPNDBF) and Open Query File (OPNQRYF). The OPNDBF command is useful in an initial program in a job for opening shared files. The OPNQRYF command is very effective in selecting and arranging records outside of your program. Then, your program can use the information supplied by the OPNQRYF command to process only the data it needs.

---

### Using the Open Database File (OPNDBF) Command

Usually, when you use the OPNDBF command, you can use the defaults for the command parameter values. In some instances you may want to specify particular values, instead of using the default values, for the following parameters:

**OPTION Parameter.** Specify the \*INP option if your application programs uses input-only processing (reading records without updating records). This allows the system to read records without trying to lock each one for possible update. Specify the \*OUT option if your application programs uses output-only processing (writing records into a file but not reading or updating existing records).

**Note:** If your program does direct output operations to active records (updating by relative record number), \*ALL must be specified instead of \*OUT. If your program does direct output operations to deleted records only, \*OUT must be specified.

**MBR Parameter.** If a member, other than the first member in the file, is to be opened, you must specify the name of the member to be opened or issue an Override with Database File (OVRDBF) command before the Open Database File (OPNDBF) command.

**Note:** You must specify a member name on the OVRDBF command to use a member (other than the first member) to open in subsequent programs.

**OPNID Parameter.** If an identifier other than the file name is to be used, you must specify it. The open identifier can be used in other CL commands to process the file. For example, the Close File (CLOF) command uses the identifier to specify which file is to be closed.

**ACCPATH Parameter.** If the file has a keyed sequence access path and either (1) the open option is \*OUT, or (2) the open option is \*INP or \*ALL, but your program does not use the keyed sequence access path, then you can specify ACCPTH(\*ARRIVAL) on the OPNDBF parameter. Ignoring the keyed sequence access path can improve your job's performance.

**SEQONLY Parameter.** Specify \*YES if subsequent application programs process the file sequentially. This parameter can also be used to specify the number of records that should be transferred between the system data buffers and the program data buffers. SEQONLY(\*YES) is not allowed unless OPTION(\*INP) or OPTION(\*OUT) is also specified on the Open Database File (OPNDBF) command. Sequential-only processing should not be used with a keyed sequence access path file unless the physical data is in access path order.

**COMMIT Parameter.** Specify \*YES if the application programs use commitment control. If you specify \*YES you must be running in a commitment control environment (the Start Commitment Control [STRCMTCTL] command was processed) or the OPNDBF command will fail. Use the default of \*NO if the application programs do not use commitment control.

**OPNSCOPE Parameter.** Specifies the scoping of the open data path (ODP). Specify \*ACTGRPDFN if the request is from the default activation group, and the ODP is to be scoped to the call level of the program issuing the command. If the request is from any other activation group, the ODP is scoped to that activation group. Specify \*ACTGRP if the ODP is to be scoped to the activation group of the program issuing the command. Specify \*JOB if the ODP is to be scoped to the job. If you specify this parameter and the TYPE parameter you get an error message.

**DUPKEYCHK Parameter.** Specify whether or not you want duplicate key feedback. If you specify \*YES, duplicate key feedback is returned on I/O operations. If you specify \*NO, duplicate key feedback is not returned on I/O operations. Use the default (\*NO) if the application programs are not written in the COBOL/400 language or C/400\* language, or if your COBOL or C programs do not use the duplicate-key feedback information that is returned.

**TYPE Parameter.** Specify what you wish to happen when exceptions that are not monitored occur in your application program. If you specify \*NORMAL one of the following can happen:

- Your program can issue a Reclaim Resources (RCLRSC) command to close the files opened at a higher level in the call stack than the program issuing the RCLRSC command.

- The high-level language you are using can perform a close operation.

Specify \*PERM if you want to continue the application without opening the files again. TYPE(\*NORMAL) causes files to be closed if both of the following occur:

- Your program receives an error message
- The files are opened at a higher level in the call stack.

TYPE(\*PERM) allows the files to remain open even if an error message is received. Do not specify this parameter if you specified the OPNSCOPE parameter.

---

## Using the Open Query File (OPNQRYF) Command

The Open Query File (OPNQRYF) command is a CL command that allows you to perform many data processing functions on database files. Essentially, the OPNQRYF command acts as a filter between the processing program and the database records. The database file can be a physical or logical file.

Unlike a database file created with the Create Physical File (CRTPF) command or the Create Logical File (CRTLFL) command, the OPNQRYF command creates only a temporary file for processing the data, it does not create a permanent file.

To understand the OPNQRYF command support, you should already be familiar with database concepts such as physical and logical files, key fields, record formats, and join logical files.

The OPNQRYF command has functions similar to those in DDS, and the CRTPF and CRTLFL commands. DDS requires source statements and a separate step to create the file. OPNQRYF allows a dynamic definition without using DDS. The OPNQRYF command does not support all of the DDS functions, but it supports significant functions that go beyond the capabilities of DDS.

In addition, Query/400 can be used to perform some of the function the OPNQRYF command performs. However, the OPNQRYF command is more useful as a programmer's tool.

The OPNQRYF command parameters also have many functions similar to the SQL/400 SELECT functions. For example, the FILE parameter is similar to the SQL/400 FROM statement, the QRYSLT parameter is similar to the SQL/400 WHERE statement, the GRPFLD parameter is similar to the SQL/400 GROUP BY statement, and the GRPSLT parameter is similar to the SQL/400 HAVING statement. For more information about the SQL/400 language, see the *SQL/400\* Programmer's Guide*.

The following is a list of the major functions supplied by OPNQRYF. Each of these functions is described later in this section.

- Dynamic record selection
- Dynamic keyed sequence access path
- Dynamic keyed sequence access path over a join
- Dynamic join
- Handling missing records in secondary join files
- Unique-key processing
- Mapped field definitions
- Group processing
- Final total-only processing

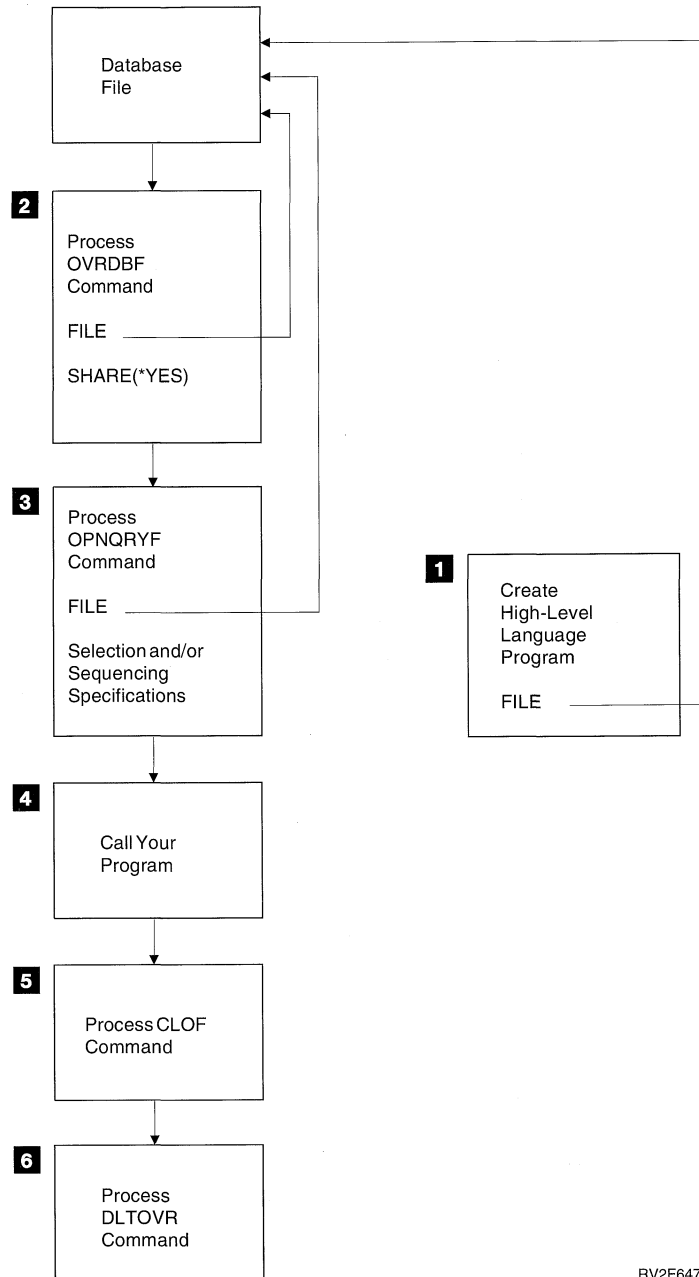
- Improving performance
- Open Query Identifier (ID)
- Sort sequence processing

To understand the OPNQRYP command, you must be familiar with its two processing approaches: using a format in the file, and using a file with a different format. The typical use of the OPNQRYP command is to select, arrange, and format the data so it can be read sequentially by your high-level language program.

See the *CL Reference* for OPNQRYP command syntax and parameter descriptions.

## Using an Existing Record Format in the File

Assume you only want your program to process the records in which the *Code* field is equal to D. You create the program as if there were only records with a D in the *Code* field. That is, you do not code any selection operations in the program. You then run the OPNQRYP command, and specify that only the records with a D in the *Code* field are to be returned to the program. The OPNQRYP command does the record selection and your program processes only the records that meet the selection values. You can use this approach to select a set of records, return records in a different sequence than they are stored, or both. The following is an example of using the OPNQRYP command to select and sequence records:



RV2F647-0

- 1** Create the high-level language program to process the database file as you would any normal program using externally described data. Only one format can be used, and it must exist in the file.
- 2** Run the OVRDBF command specifying the file and member to be processed and SHARE(\*YES). (If the member is permanently changed to SHARE(\*YES) and the first or only member is the one you want to use, this step is not necessary.)

The OVRDBF command can be run after the OPNQRYP command, unless you want to override the file name specified in the OPNQRYP command. In this discussion and in the examples, the OVRDBF command is shown first.

Some restrictions are placed on using the OVRDBF command with the OPNQRYP command. For example, MBR(\*ALL) causes an error message

and the file is not opened. Refer to “Considerations for Files Shared in a Job” on page 6-53 for more information.

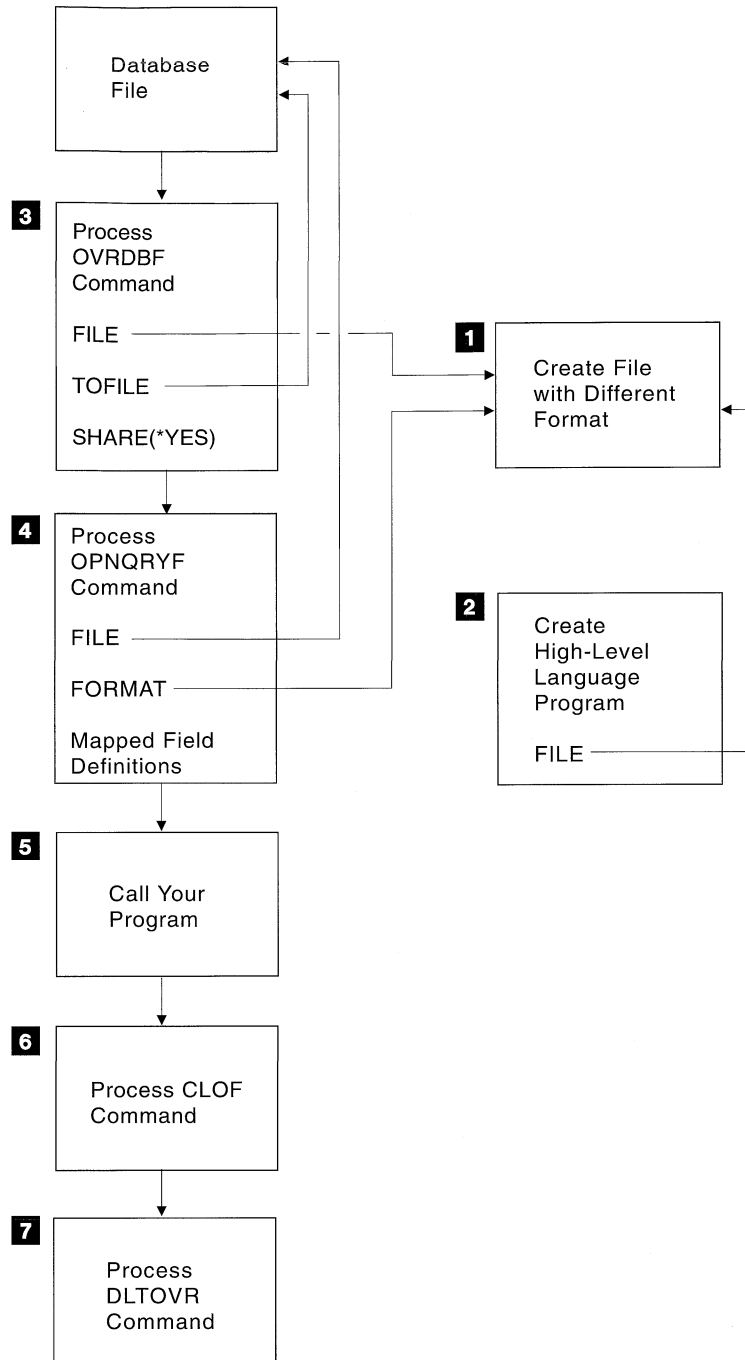
- 3** Run the OPNQRYF command, specifying the database file, member, format names, any selection options, any sequencing options, and the scope of influence for the opened file.
- 4** Call the high-level language program you created in step 1. Besides using a high-level language, the Copy from Query File (CPYFRMQRYP) command can also be used to process the file created by the OPNQRYF command. Other CL commands (for example, the Copy File [CPYF] and the Display Physical File Member [DSPPFM] commands) and utilities (for example, Query) do not work on files created with the OPNQRYF command.
- 5** Close the file that you opened in step 3, unless you want the file to remain open. The Close File (CLOF) command can be used to close the file.
- 6** Delete the override specified in step 2 with the Delete Override (DLTOVR) command. It may not always be necessary to delete the override, but the command is shown in all the examples for consistency.

## Using a File with a Different Record Format

For more advanced functions of the Open Query File (OPNQRYF) command (such as dynamically joining records from different files), you must define a new file that contains a different record format. This new file is a separate file from the one you are going to process. This new file contains the fields that you want to create with the OPNQRYF command. This powerful capability also lets you define fields that do not currently exist in your database records, but can be derived from them.

When you code your high-level language program, specify the name of the file with the different format so the externally described field definitions of both existing and derived fields can be processed by the program.

Before calling your high-level language program, you must specify an Override with Database File (OVRDBF) command to direct your program file name to the open query file. On the OPNQRYF command, specify both the database file and the new file with the special format to be used by your high-level language program. If the file you are querying does not have SHARE(\*YES) specified, you must specify SHARE(\*YES) on the OVRDBF command.



RSLH299-4

- 1** Specify the DDS for the file with the different record format, and create the file. This file contains the fields that you want to process with your high-level language program. Normally, data is not contained in this file, and it does not require a member. You normally create this file as a physical file without keys. A field reference file can be used to describe the fields. The record format name can be different from the record format name in the database file that is specified. You can use any database or DDM file for this function. The file could be a logical file and it could be indexed. It could have one or more members, with or without data.

- 2** Create the high-level language program to process the file with the record format that you created in step 1. In this program, do not name the database file that contains the data.
- 3** Run the Override with Database File (OVRDBF) command. Specify the name of the file with the different (new) record format on the FILE parameter. Specify the name of the database file that you want to query on the TOFILE parameter. You can also specify a member name on the MBR parameter. If the database member you are querying does not have SHARE(\*YES) specified, you must also specify SHARE(\*YES) on the OVRDBF command.
- 4** Run the Open Query File (OPNQRYF) command. Specify the database file to be queried on the FILE parameter, and specify the name of the file with the different (new) format that was created in step 1 on the FORMAT parameter. Mapped field definitions can be required on the OPNQRYF command to describe how to map the data from the database file into the format that was created in step 1. You can also specify selection options, sequencing options, and the scope of influence for the opened file.
- 5** Call the high-level language program you created in step 2.
- 6** The first file named in step 4 for the FILE parameter was opened with OPNQRYF as SHARE(\*YES) and is still open. The file must be closed. The Close File (CLOF) command can be used.
- 7** Delete the override that was specified in step 3.

The previous steps show the normal flow using externally described data. It is not necessary to create unique DDS and record formats for each OPNQRYF command. You can reuse an existing record format. However, all fields in the record format must be actual fields in the real database file or defined by mapped field definitions. If you use program-described data, you can create the program at any time.

You can use the file created in step 1 to hold the data created by the Open Query File (OPNQRYF) command. For example, you can replace step 5 with a high-level language processing program that copies data to the file with the different format, or you may use the Copy from Query File (CPYFRMQRYP) command. The Copy File (CPYF) command cannot be used. You can then follow step 5 with the CPYF command or Query.

## OPNQRYF Examples

The following sections describe how to specify both the OPNQRYF parameters for each of the major functions discussed earlier and how to use the Open Query File command with your high-level language program.

### Notes:

1. If you run the OPNQRYF command from a command entry line with the OPNSCOPE(\*ACTGRPDFN) or TYPE(\*NORMAL) parameter option, error messages that occur after the OPNQRYF command successfully runs will not close the file. Such messages would have closed the file prior to Version 2 Release 3 when TYPE(\*NORMAL) was used. The system automatically runs the Reclaim Resources (RCLRSC) command if an error message occurs, except for message CPF0001, which is sent when the system detects an error in the command. However, the RCLRSC command only closes files opened from the default activation group at a higher level in the call stack than the level at which the RCLRSC command was run.



2. After running a program that uses the Open Query File command for sequential processing, the file position is normally at the end of the file. If you want to run the same program or a different program with the same files, you must position the file or close the file and open it with the same OPNQRYF command. You can position the file with the Position Database File (POSDBF) command. In some cases, a high-level language program statement can be used.

## The Zero Length Literal and the Contains (\*CT) Function

The concept of a *zero length literal* was introduced in Version 2, Release 1, Modification 1. In the OPNQRYF command, a zero length literal is denoted as a quoted string with nothing, not even a blank, between the quotes ("").

Zero length literal support changes the results of a comparison when used as the compare argument of the contains (\*CT) function. Consider the statement:

```
QRYSLT('field *CT ""')
```

With zero length literal support, the statement returns records that contain anything. It is, in essence, a wildcard comparison for any number of characters followed by any number of characters. It is equivalent to:

```
'field = %WLDCRD("**")'
```

Before zero length literal support, (before Version 2, Release 1, Modification 1), the argument (") was interpreted as a single-byte blank. The statement returned records that contained a single blank somewhere in the field. It was, in essence, a wildcard comparison for any number of characters, followed by a blank, followed by any number of characters. It was equivalent to:

```
'field = %WLDCRD("* *")'
```

### To Repeat Pre-Zero Length Literal Support Results

To get pre-Version 2, Release 1, Modification 1 results with the contains function, you must code the QRYSLT to explicitly look for the blank:

```
QRYSLT('field *CT " "')
```

## Selecting Records without Using DDS

Dynamic record selection allows you to request a subset of the records in a file without using DDS. For example, you can select records that have a specific value or range of values (for example, all customer numbers between 1000 and 1050). The Open Query File (OPNQRYF) command allows you to combine these and other selection functions to produce powerful record selection capabilities.

### Examples of Selecting Records Using the Open Query File (OPNQRYF) Command

In all of the following examples, it is assumed that a single-format database file (physical or logical) is being processed. (The FILE parameter on the OPNQRYF command allows you to specify a record format name if the file is a multiple format logical file.)

See the OPNQRYF command in the *CL Reference* for a complete description of the format of expressions used with the QRYSLT parameter.

**Example 1:** Selecting records with a specific value

Assume you want to select all the records from FILEA where the value of the *Code* field is D. Your processing program is PGMB. PGMB only sees the records that meet the selection value (you do not have to test in your program).

**Note:** You can specify parameters easier by using the prompt function for the OPNQRYF command. For example, you can specify an expression for the QRYSLT parameter without the surrounding delimiters because the system will add the apostrophes.

Specify the following:

```
OVRDBF      FILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA) QRYSLT('CODE *EQ "D" ')
CALL        PGM(PGMB)
CLOF        OPNID(FILEA)
DLTOVR      FILE(FILEA)
```

**Notes:**

1. The entire expression in the QRYSLT parameter must be enclosed in apostrophes.
2. When specifying field names in the OPNQRYF command, the names in the record are not enclosed in apostrophes.
3. Character literals must be enclosed by quotation marks or two apostrophes. (The quotation mark character is used in the examples.) It is important to place the character(s) between the quotation marks in either uppercase or lowercase to match the value you want to find in the database. (The examples are all shown in uppercase.)
4. To request a selection against a numeric constant, specify:

```
OPNQRYF     FILE(FILEA) QRYSLT('AMT *GT 1000.00')
```

Notice that numeric constants are *not* enclosed by two apostrophes (quotation marks).

5. When comparing a field value to a CL variable, use apostrophes as follows (only character CL variables can be used):

- If doing selection against a character, date, time, or timestamp field, specify:

```
OPNQRYF     FILE(FILEA) QRYSLT('"' *CAT &CHAR *CAT "' *EQ FIELDA')
```

or, in reverse order:

```
OPNQRYF     FILE(FILEA) QRYSLT('FIELDA *EQ "' *CAT &CHAR *CAT "'')
```

Notice that apostrophes and quotation marks enclose the CL variables and \*CAT operators.

- If doing selection against a numeric field, specify:

```
OPNQRYF     FILE(FILEA) QRYSLT(&CHARNUM *CAT ' *EQ NUM')
```

or, in reverse order:

```
OPNQRYF     FILE(FILEA) QRYSLT('NUM *EQ ' *CAT &CHARNUM)
```

Notice that apostrophes enclose the field and operator only.

When comparing two fields or constants, the data types must be compatible. The following table describes the valid comparisons.

Figure 6-1. Valid Data Type Comparisons for the OPNQRYF Command

	Any Numeric	Character	Date <sup>1</sup>	Time <sup>1</sup>	Timestamp <sup>1</sup>
<b>Any Numeric</b>	Valid	Not Valid	Not Valid	Not Valid	Not Valid
<b>Character</b>	Not Valid	Valid	Valid <sup>2</sup>	Valid <sup>2</sup>	Valid <sup>2</sup>
<b>Date<sup>1</sup></b>	Not Valid	Valid <sup>2</sup>	Valid	Not Valid	Not Valid
<b>Time<sup>1</sup></b>	Not Valid	Valid <sup>2</sup>	Not Valid	Valid	Not Valid
<b>Timestamp<sup>1</sup></b>	Not Valid	Valid <sup>2</sup>	Not Valid	Not Valid	Valid

<sup>1</sup> Date, time, and timestamp data types can be represented by fields and expressions, but not constants; however, character constants can represent date, time, or timestamp values.

<sup>2</sup> The character field or constant must represent a valid date value if compared to a date data type, a valid time value if compared to a time data type, or a valid timestamp value if compared to a timestamp data type.

**Note:** For DBCS information, see Appendix B, “Double-Byte Character Set (DBCS) Considerations.”

The performance of record selection can be greatly enhanced if some file on the system uses the field being selected in a keyed sequence access path. This allows the system to quickly access only the records that meet the selection values. If no such access path exists, the system must read every record to determine if it meets the selection values.

Even if an access path exists on the field you want to select from, the system may not use the access path. For example, if it is faster for the system to process the data in arrival sequence, it will do so. See the discussion in “Performance Considerations” on page 6-45 for more details.

**Example 2:** Selecting records with a specific date value

Assume you want to process all records in which the *Date* field in the record is the same as the current date. Also assume the *Date* field is in the same format as the system date. In a CL program, you can specify:

```
DCL          VAR(&CURDAT) TYPE(*CHAR) LEN(6)
RTVSYVAL    SYSVAL(QDATE) RTNVAR(&CURDAT)
OVRDBF      FILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA) QRYSLT('"' *CAT &CURDAT *CAT '"' *EQ DATE')
CALL        PGM(PGMB)
CLOF        OPNID(FILEA)
DLTOVR      FILE(FILEA)
```

A CL variable is assigned with a leading ampersand (&) and is not enclosed in apostrophes. The quotation mark is enclosed in apostrophes as is the string \*EQ DATE. The quotation mark is concatenated to the value of the CL variable, as is the remainder of the QRYSLT string.

It is important to know whether the data in the database is defined as character, date, time, timestamp, or numeric. In the preceding example, the *Date* field is assumed to be character.

If the *DATE* field is defined as date data type, the preceding example could be specified as:

```
OVRDBF FILE(FILEA) SHARE(*YES)
OPNQRYF FILE(FILEA) QRYSLT('%CURDATE *EQ DATE')
CALL PGM(PGMB)
CLOF OPENID(FILEA)
DLTOVR FILE(FILEA)
```

**Note:** The date field does not have to have the same format as the system date.

You could also specify the example as:

```
DCL VAR(&CVTDAT) TYPE(*CHAR) LEN(6)
DCL VAR(&CURDAT) TYPE(*CHAR) LEN(8)
RTVSYSVAL SYSVAL(QDATE) RTNVAR(&CVTDAT)
CVTDAT DATE(&CVTDAT) TOVAR(&CURDAT) TOSEP(/)
OVRDBF FILE(FILEA) SHARE(*YES)
OPNQRYF FILE(FILEA)
      QRYSLT('"' *CAT &CURDAT *CAT '"' *EQ DATE')
CALL PGM(PGMB)
CLOF OPNID (FILEA)
DLTOVR FILE(FILEA)
```

This is where *DATE* has a date data type in FILEA, the job default date format is MMDDYY, and the job default date separator is the slash (/).

**Note:** For any character representation of a date in one of the following formats, MMDDYY, DDMMYY, YYMMDD, or Julian, the job default date format and separator must be the same to be recognized.

If, instead, you were using a constant, the QRYSLT would be specified as follows:

```
QRYSLT('"12/31/87" *EQ DATE')
```

The job default date format must be MMDDYY and the job default separator must be the slash (/).

**Note:** For any character representation of a date in one of the following formats, MMDDYY, DDMMYY, YYMMDD, or Julian, the job default date format and separator must be the same to be recognized.

If a numeric field exists in the database and you want to compare it to a variable, only a character variable can be used. For example, to select all records where a packed *Date* field is greater than a variable, you must ensure the variable is in character form. Normally, this will mean that before the Open Query File (OPNQRYF) command, you use the Change Variable (CHGVAR) command to change the variable from a decimal field to a character field. The CHGVAR command would be specified as follows:

```
CHGVAR VAR(&CHARVAR) VALUE('123188')
```

The QRYSLT parameter would be specified as follows (see the difference from the preceding examples):

```
QRYSLT(&CHARVAR *CAT ' *GT DATE')
```

If, instead, you were using a constant, the QRYSLT statement would be specified as follows:

```
QRYSLT('123187 *GT DATE')
```

**Example 3:** Selecting records in a range of values

Assume you have a *Date* field specified in the character format YYMMDD and with the “.” separator, and you want to process all records for 1988. You can specify:

```
OVRDBF      FILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA) QRYSLT('DATE *EQ %RANGE("88.01.01" +
                                "88.12.31") ')
CALL        PGM(PGMC)
CLOF        OPNID(FILEA)
DLTOVR      FILE(FILEA)
```

This example would also work if the *DATE* field has a date data type, the job default date format is YYMMDD, and the job default date separator is the period (.).

**Note:** For any character representation of a date in one of the following formats, MMDDYY, DDMYY, YYMMDD, or Julian, the job default date format and separator must be the same to be recognized.

If the ranges are variables defined as character data types, and the *DATE* field is defined as a character data type, specify the *QRYSLT* parameter as follows:

```
QRYSLT('DATE *EQ %RANGE(" *CAT &LORNG *CAT "' *BCAT "' +
        *CAT &HIRNG *CAT "')')
```

However, if the *DATE* field is defined as a numeric data type, specify the *QRYSLT* parameter as follows:

```
QRYSLT('DATE *EQ %RANGE(' *CAT &LORNG *BCAT &HIRNG *CAT ')')
```

**Note:** \*BCAT can be used if the *QRYSLT* parameter is in a CL program, but it is not allowed in an interactive command.

**Example 4:** Selecting records using the contains function

Assume you want to process all records in which the *Addr* field contains the street named BROADWAY. The contains (\*CT) function determines if the characters appear anywhere in the named field. You can specify:

```
OVRDBF      FILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA) QRYSLT('ADDR *CT "BROADWAY" ')
CALL        PGM(PGMC)
CLOF        OPNID(FILEA)
DLTOVR      FILE(FILEA)
```

In this example, assume that the data is in uppercase in the database record. If the data was in lowercase or mixed case, you could specify a translation function to translate the lowercase or mixed case data to uppercase before the comparison is made. The system-provided table QSYSTRNTBL translates the letters a through z to uppercase. (You could use any translation table to perform the translation.)

Therefore, you can specify:

```
OVRDBF      FILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA) QRYSLT('%XLATE(ADDR QSYSTRNTBL) *CT +
                                "BROADWAY" ')
CALL        PGM(PGMC)
CLOF        OPNID(FILEA)
DLTOVR      FILE(FILEA)
```

When the %XLATE function is used on the QRYSLT statement, the value of the field passed to the high-level language program appears as it is in the database. You can force the field to appear in uppercase using the %XLATE function on the MAPFLD parameter.

**Example 5:** Selecting records using multiple fields

Assume you want to process all records in which either the *Amt* field is equal to zero, or the *Lstdat* field (YYMMDD order in character format) is equal to or less than 88-12-31. You can specify:

```
OVRDBF      FILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA) QRYSLT('AMT *EQ 0 *OR LSTDAT +
                    *LE "88-12-31" ')
CALL        PGM(PGMC)
CLOF        OPNID(FILEA)
DLTOVR      FILE(FILEA)
```

This example would also work if the *LSTDAT* field has a date data type. The *LSTDAT* field may be in any valid date format; however, the job default date format must be YYMMDD and the job default date separator must be the dash (-).

**Note:** For any character representation of a date in one of the following formats, MMDDYY, DDMMYY, YYMMDD, or Julian, the job default date format and separator must be the same to be recognized.

If variables are used, the QRYSLT parameter is typed as follows:

```
QRYSLT('AMT *EQ ' *CAT &VARAMT *CAT ' *OR +
        LSTDAT *LE '' *CAT &VARDAT *CAT ''')
```

or, typed in reverse order:

```
QRYSLT(''' *CAT &VARDAT *CAT '' *GT LSTDAT *OR ' +
        *CAT &VARAMT *CAT ' *EQ AMT')
```

Note that the &VARAMT variable must be defined as a character type. If the variable is passed to your CL program as a numeric type, you must convert it to a character type to allow concatenation. You can use the Change Variable (CHGVAR) command to do this conversion.

**Example 6:** Using the Open Query File (OPNQRYF) command many times in a program

You can use the OPNQRYF command more than once in a high-level language program. For example, assume you want to prompt the user for some selection values, then display one or more pages of records. At the end of the first request for records, the user may want to specify other selection values and display those records. This can be done by doing the following:

1. Before calling the high-level language program, use an Override with Database File (OVRDBF) command to specify SHARE(\*YES).
2. In the high-level language program, prompt the user for the selection values.
3. Pass the selection values to a CL program that issues the OPNQRYF command (or run the command with a call to program QCMDXC). The file must be closed before your program processes the OPNQRYF command. You normally use the Close File (CLOF) command and monitor for the file not being open.

4. Return to the high-level language program.
5. Open the file in the high-level language program.
6. Process the records.
7. Close the file in the program.
8. Return to step 2.

When the program completes, run the Close File (CLOF) command or the Reclaim Resources (RCLRSC) command to close the file, then delete the Override with Database File command specified in step 1.

**Note:** An override command in a called CL program does not affect the open in the main program. All overrides are implicitly deleted when the program is ended. (However, you can use a call to program QCMDEXC from your high-level language program to specify an override, if needed.)

**Example 7:** Mapping fields for packed numeric data fields

Assume you have a packed decimal *Date* field in the format MMDDYY and you want to select all the records for the year 1988. You cannot select records directly from a portion of a packed decimal field, but you can use the MAPFLD parameter on the OPNQRYF command to create a new field that you can then use for selecting part of the field.

The format of each mapped field definition is:

(result field 'expression' attributes)

where:

- |              |   |                                                                                                                                                                                                               |
|--------------|---|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| result field | = | The name of the result field.                                                                                                                                                                                 |
| expression   | = | How the result field should be derived. The expression can include substring, other built-in functions, or mathematical statements.                                                                           |
| attributes   | = | The optional attributes of the result field. If no attributes are given (or the field is not defined in a file), the OPNQRYF command calculates a field attribute determined by the fields in the expression. |

```
OVRDBF      FILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA) QRYSLT('YEAR *EQ "88" ') +
            MAPFLD((CHAR6 '%DIGITS(DATE)') +
            (YEAR '%SST(CHAR6 5 2)') *CHAR 2))
CALL        PGM(PGMC)
CLOF        OPNID(FILEA)
DLTOVR      FILE(FILEA)
```

In this example, if DATE was a date data type, it could be specified as follows:

```
OPNQRYF FILE(FILEA) +
QRYSLT ('YEAR *EQ 88') +
MAPFLD((YEAR '%YEAR(DATE)'))
```

The first mapped field definition specifies that the *Char6* field be created from the packed decimal *Date* field. The %DIGITS function converts from packed decimal to character and ignores any decimal definitions (that is, 1234.56 is converted to '123456'). Because no definition of the *Char6* field is specified, the system assigns

a length of 6. The second mapped field defines the *Year* field as type \*CHAR (character) and length 2. The expression uses the substring function to map the last 2 characters of the *Char6* field into the *Year* field.

Note that the mapped field definitions are processed in the order in which they are specified. In this example, the *Date* field was converted to character and assigned to the *Char6* field. Then, the last two digits of the *Char6* field (the year) were assigned to the *Year* field. Any changes to this order would have produced an incorrect result.

**Note:** Mapped field definitions are always processed before the QRYSLT parameter is evaluated.

You could accomplish the same result by specifying the substring on the QRYSLT parameter and dropping one of the mapped field definitions as follows:

```
OPNQRYF      FILE(FILEA) +
              QRYSLT('%SST(CHAR6 5 2) *EQ "88" ') +
              MAPFLD((CHAR6 '%DIGITS(DATE)'))
```

**Example 8:** Using the “wildcard” function

Assume you have a packed decimal *Date* field in the format MMDDYY and you want to select the records for March 1988. To do this, you can specify:

```
OVRDBF      FILE(FILEA) SHARE(*YES)
OPNQRYF      FILE(FILEA) +
              QRYSLT('%DIGITS(DATE) *EQ %WLDCRD("03__88"')
CALL         PGM(PGMC)
CLOF        OPNID(FILEA)
DLTOVR      FILE(FILEA)
```

Note that the only time the MAPFLD parameter is needed to define a database field for the result of the %DIGITS function is when the result needs to be used with a function that only supports a simple field name (not a function or expression) as an argument. The %WLDCRD operation has no such restriction on the operand that appears before the \*EQ operator.

Note that although the field in the database is in numeric form, double apostrophes surround the literal to make its definition the same as the *Char6* field. The wildcard function is not supported for DATE, TIME, or TIMESTAMP data types.

The %WLDCRD function lets you select any records that match your selection values, in which the underline ( ) will match any single character value. The two underline characters in Example 8 allow any day in the month of March to be selected. The %WLDCRD function also allows you to name the wild card character (underline is the default).

The wild card function supports two different forms:

- A fixed-position wild card as shown in the previous example in which the underline (or your designated character) matches any single character as in the following example:

```
QRYSLT('FLDA *EQ %WLDCRD("A_C"')
```

This compares successfully to ABC, ACC, ADC, AxC, and so on. In this example, the field being analyzed only compares correctly if it is exactly 3 char-



acters in length. If the field is longer than 3 characters, you also need the second form of wild card support.

- A variable-position wild card will match any zero or more characters. The Open Query File (OPNQUERY) command uses an asterisk (\*) for this type of wild card variable character or you can specify your own character. An asterisk is used in the following example:

```
QRYSLT('FLDB *EQ %WLDCRD("A*C*")')
```

This compares successfully to AC, ABC, AxC, ABCD, AxxxxxxC, and so on. The asterisk causes the command to ignore any intervening characters if they exist. Notice that in this example the asterisk is specified both before and after the character or characters that can appear later in the field. If the asterisk were omitted from the end of the search argument, it causes a selection only if the field ends with the character C.

You must specify an asterisk at the start of the wild card string if you want to select records where the remainder of the pattern starts anywhere in the field. Similarly, the pattern string must end with an asterisk if you want to select records where the remainder of the pattern ends anywhere in the field.

For example, you can specify:

```
QRYSLT('FLDB *EQ %WLDCRD("*ABC*DEF*")')
```

You get a match on ABCDEF, ABCxDEF, ABCxDEFx, ABCxxxxxxDEF, ABCxxxDEFxxx, xABCDEF, xABCxDEFx, and so on.

You can combine the two wildcard functions as in the following example:

```
QRYSLT('FLDB *EQ %WLDCRD("ABC_*DEF*")')
```

You get a match on ABCxDEF, ABCxxxxxxDEF, ABCxxxDEFxxx, and so on. The underline forces at least one character to appear between the ABC and DEF (for example, ABCDEF would not match).

Assume you have a *Name* field that contains:

```
JOHNS  
JOHNS SMITH  
JOHNSON  
JOHNSTON
```

If you specify the following you will only get the first record:

```
QRYSLT('NAME *EQ "JOHNS"')
```

You would not select the other records because a comparison is made with blanks added to the value you specified. The way to select all four names is to specify:

```
QRYSLT('NAME *EQ %WLDCRD("JOHNS*")')
```

**Note:** For information about using the %WLDCRD function for DBCS, see Appendix B, "Double-Byte Character Set (DBCS) Considerations."

### **Example 9:** Using complex selection statements

Complex selection statements can also be specified. For example, you can specify:

```
QRYSLT('DATE *EQ "880101" *AND AMT *GT 5000.00')
```

```
QRYSLT('DATE *EQ "880101" *OR AMT *GT 5000.00')
```

You can also specify:

```
QRYSLT('CODE *EQ "A" *AND TYPE *EQ "X" *OR CODE *EQ "B")
```

The rules governing the priority of processing the operators are described in the *CL Reference* manual. Some of the rules are:

- The \*AND operations are processed first; therefore, the record would be selected if:

The *Code* field = "A" and The *Type* field = "X"

or

The *Code* field = "B"

- Parentheses can be used to control how the expression is handled, as in the following example:

```
QRYSLT('(CODE *EQ "A" *OR CODE *EQ "B") *AND TYPE *EQ "X" +  
      *OR CODE *EQ "C"')
```

The *Code* field = "A" and The *Type* field= "X"

or

The *Code* field = "B" and The *Type* field = "X"

or

The *Code* field = "C"

You can also use the symbols described in the *CL Reference* manual instead of the abbreviated form (for example, you can use = instead of \*EQ) as in the following example:

```
QRYSLT('CODE = "A" & TYPE = "X" | AMT > 5000.00')
```

This command selects all records in which:

The *Code* field = "A" and The *Type* field = "X"

or

The *Amt* field > 5000.00

A complex selection statement can also be written, as in the following example:

```
QRYSLT('CUSNBR = %RANGE("60000" "69999") & TYPE = "B" +  
      & SALES>0 & ACCRCV / SALES>.3')
```

This command selects all records in which:

The *Cusnbr* field is in the range 60000-69999 and

The *Type* field = "B" and

The *Sales* fields are greater than 0 and

*Accrcv* divided by *Sales* is greater than 30 percent

**Example 10:** Using coded character set identifiers (CCSIDs)

For general information about CCSIDs, see the *National Language Support Planning Guide*.

Each character and DBCS field in all database files is tagged with a CCSID. This CCSID allows you to further define the data stored in the file so that any comparison, join, or display of the fields is performed in a meaningful way. For example, if you compared FIELD1 in FILE1 where FIELD1 has a CCSID of 37 (USA) to FIELD2 in FILE2 where FIELD2 has a CCSID of 273 (Austria, Germany) appropriate mapping would occur to make the comparison meaningful.

```
OPNQRYF FILE(FILEA FILEB) FORMAT(RESULTF) +  
        JFLD((FILEA/NAME FILEB/CUSTOMER))
```

If field NAME has a CCSID of 37 and field CUSTOMER has a CCSID of 273, the mapping of either NAME or CUSTOMER is performed during processing of the OPNQRYF command so that the join of the two fields provides a meaningful result.

Normally, constants defined in the MAPFLD, QRYSLT, and GRPSLT parameters are tagged with the CCSID defined to the current job. This suggests that when two users with different job CCSIDs run the same OPNQRYF command (or a program containing an OPNQRYF command) and the OPNQRYF has constants defined in it, the users can get different results because the CCSID tagged to the constants may cause the constants to be treated differently.

You can tag a constant with a specific CCSID by using the MAPFLD parameter. By specifying a MAPFLD whose definition consists solely of a constant and then specifying a CCSID for the MAPFLD the constant becomes tagged with the CCSID specified in the MAPFLD parameter. For example:

```
OPNQRYF FILE(FILEA) FORMAT(RESULTF) QRYSLT('NAME *EQ MAP1') +  
        MAPFLD((MAP1 '"Smith"' *CHAR 5 *N 37))
```

The constant "Smith" is tagged with the CCSID 37 regardless of the job CCSID of the user issuing the OPNQRYF command. In this example, all users would get the same result records (although the result records would be mapped to the user's job CCSID). Conversely, if the query is specified as:

```
OPNQRYF FILE(FILEA) FORMAT(RESULTF) QRYSLT('NAME *EQ "Smith"')
```

the results of the query may differ, depending on the job CCSID of the user issuing the OPNQRYF command.

| **Example 11:** Using Sort Sequence and Language Identifier

| To see how to use a sort sequence, run the examples in this section against the  
| STAFF file shown in Figure 6-2.

Figure 6-2. The STAFF File

ID	NAME	DEPT	JOB	YEARS	SALARY	COMM
10	Sanders	20	Mgr	7	18357.50	0
20	Pernal	20	Sales	8	18171.25	612.45
30	Merenghi	38	MGR	5	17506.75	0
40	OBrien	38	Sales	6	18006.00	846.55
50	Hanes	15	Mgr	10	20659.80	0
60	Quigley	38	SALES	00	16808.30	650.25
70	Rothman	15	Sales	7	16502.83	1152.00
80	James	20	Clerk	0	13504.60	128.20
90	Koonitz	42	sales	6	18001.75	1386.70
100	Plotz	42	mgr	6	18352.80	0

In the examples, the results are shown for a particular statement using each of the following:

- \*HEX sort sequence.
- Shared-weight sort sequence for language identifier ENU.
- Unique-weight sort sequence for language identifier ENU.

**Note:** ENU is chosen as a language identifier by specifying either SRTSEQ(\*LANGIDUNQ) or SRTSEQ(\*LANGIDSHR), and LANGID(ENU) in the OPNQRYF command.

The following command selects records with the value MGR in the JOB field:

```
OPNQRYF FILE(STAFF) QRYSLT('JOB *EQ "MGR"')
```

Figure 6-3 shows the record selection with the \*HEX sort sequence. The records that match the record selection criteria for the JOB field are selected exactly as specified in the QRYSLT statement; only the uppercase MGR is selected.

Figure 6-3. Using the \*HEX Sort Sequence. OPNQRYF FILE(STAFF) QRYSLT('JOB \*EQ "MGR"') SRTSEQ(\*HEX)

ID	NAME	DEPT	JOB	YEARS	SALARY	COMM
30	Merenghi	38	MGR	5	17506.75	0

Figure 6-4 shows the record selection with the shared-weight sort sequence. The records that match the record selection criteria for the JOB field are selected by treating uppercase and lowercase letters the same. With this sort sequence, mgr, Mgr, and MGR values are selected.

Figure 6-4. Using the Shared-Weight Sort Sequence. OPNQRYF FILE(STAFF) QRYSLT('JOB \*EQ "MGR"') SRTSEQ(LANGIDSHR) LANGID(ENU)

ID	NAME	DEPT	JOB	YEARS	SALARY	COMM
10	Sanders	20	Mgr	7	18357.50	0
30	Merenghi	38	MGR	5	17506.75	0
50	Hanes	15	Mgr	10	20659.80	0
100	Plotz	42	mgr	6	18352.80	0

Figure 6-5 on page 6-21 shows the record selection with the unique-weight sort sequence. The records that match the record selection criteria for the JOB field are selected by treating uppercase and lowercase letters as unique. With this sort sequence, the mgr, Mgr, and MGR values are all different. The MGR value is selected.

Figure 6-5. Using the Unique-Weight Sort Sequence. `OPNQRYF FILE(STAFF) QRYSLT('JOB *EQ "MGR") SRTSEQ(LANGIDUNQ) LANGID(ENU)`

ID	NAME	DEPT	JOB	YEARS	SALARY	COMM
30	Merenghi	38	MGR	5	17506.75	0

## Specifying a Keyed Sequence Access Path without Using DDS

The dynamic access path function allows you to specify a keyed access path for the data to be processed. If an access path already exists that can be shared, the system can share it. If a new access path is required, it is built before any records are passed to the program.

**Example 1:** Arranging records using one key field

Assume you want to process the records in FILEA arranged by the value in the *Cust* field with program PGMD. You can specify:

```
OVRDBF    FILE(FILEA) SHARE(*YES)
OPNQRYF   FILE(FILEA) KEYFLD(CUST)
CALL      PGM(PGMD)
CLOF      OPNID(FILEA)
DLTOVR    FILE(FILEA)
```

**Note:** The FORMAT parameter on the Open Query File (OPNQRYF) command is not needed because PGMD is created by specifying FILEA as the processed file. FILEA can be an arrival sequence or a keyed sequence file. If FILEA is keyed, its key field can be the *Cust* field or a totally different field.

**Example 2:** Arranging records using multiple key fields

If you want the records to be processed by *Cust* sequence and then by *Date* in *Cust*, specify:

```
OPNQRYF   FILE(FILEA) KEYFLD(CUST DATE)
```

If you want the *Date* to appear in descending sequence, specify:

```
OPNQRYF   FILE(FILEA) KEYFLD((CUST) (DATE *DESCEND))
```

In these two examples, the FORMAT parameter is not used. (If a different format is defined, all key fields must exist in the format.)

**Example 3:** Arranging records using a unique-weight sort sequence.

To process the records by the JOB field values with a unique-weight sort sequence using the STAFF file in Figure 6-2 on page 6-19, specify:

```
OPNQRYF FILE(STAFF) KEYFLD(JOB) SRTSEQ(*LANGIDUNQ) LANGID(ENU)
```

This query results in a JOB field in the following sequence:

Clerk

```

|           mgr
|           Mgr
|           Mgr
|           MGR
|           sales
|           Sales
|           Sales
|           Sales
|           SALES

```

**Example 4:** Arranging records using a shared-weight sort sequence.

To process the records by the JOB field values with a unique-weight sort sequence using the STAFF file in Figure 6-2 on page 6-19, specify:

```
OPNQRYF FILE(STAFF) KEYFLD(JOB) SRTSEQ(*LANGIDSHR) LANGID(ENU)
```

The results from this query will be similar to the results in Example 3. The *mgr* and *sales* entries could be in any sequence because the uppercase and lowercase letters are treated as equals. That is, the shared-weight sort sequence treats mgr, Mgr, and MGR as equal values. Likewise, sales, Sales, and SALES are treated as equal values.

## Specifying Key Fields from Different Files

A dynamic keyed sequence access path over a join logical file allows you to specify a processing sequence in which the keys can be in different physical files (DDS restricts the keys to the primary file).

The specification is identical to the previous method. The access path is specified using whatever key fields are required. There is no restriction on which physical file the key fields are in. However, if a key field exists in other than the primary file of a join specification, the system must make a temporary copy of the joined records. The system must also build a keyed sequence access path over the copied records before the query file is opened. The key fields must exist in the format identified on the FORMAT parameter.

**Example 1:** Using a field in a secondary file as a key field

Assume you already have a join logical file named JOINLF. FILEX is specified as the primary file and is joined to FILEY. You want to process the records in JOINLF by the *Descrp* field which is in FILEY.

Assume the file record formats contain the following fields:

FILEX	FILEY	JOINLF
Item	Item	Item
Qty	Descrp	Qty
		Descrp

You can specify:

```

OVRDBF    FILE(JOINLF) SHARE(*YES)
OPNQRYF   FILE(JOINLF) KEYFLD(DESCRP)
CALL      PGM(PGMC)
CLOF      OPNID(JOINLF)
DLTOVR    FILE(JOINLF)

```

If you want to arrange the records by *Qty* in *Descrp* (*Descrp* is the primary key field and *Qty* is a secondary key field) you can specify:

```
OPNQRYF   FILE(JOINLF) KEYFLD(DESCRP QTY)
```

## Dynamically Joining Database Files without DDS

The dynamic join function allows you to join files without having to first specify DDS and create a join logical file. You must use the `FORMAT` parameter on the Open Query File (`OPNQRYF`) command to specify the record format for the join. You can join any physical or logical file including a join logical file and a view (DDS does not allow you to join logical files). You can specify either a keyed or arrival sequence access path. If keys are specified, they can be from any of the files included in the join (DDS restricts keys to just the primary file).

In the following examples, it is assumed that the file specified on the `FORMAT` parameter was created. You will normally want to create the file before you create the processing program so you can use the externally described data definitions.

The default for the join order (`JORDER`) parameter is used in all of the following examples. The default for the `JORDER` parameter is `*ANY`, which tells the system that it can determine the order in which to join the files. That is, the system determines which file to use as the primary file and which as the secondary files. This allows the system to try to improve the performance of the join function.

The join criterion, like the record selection criterion, is affected by the sort sequence (`SRTSEQ`) and the language identifier (`LANGID`) specified (see "Example 11" on page 6-19).

### **Example 1:** Dynamically joining files

Assume you want to join `FILEA` and `FILEB`. Assume the files contain the following fields:

FILEA	FILEB	JOINAB
Cust	Cust	Cust
Name	Amt	Name
Addr		Amt

The join field is `Cust` which exists in both files. Any record format name can be specified in the Open Query File (`OPNQRYF`) command for the join file. The file does not need a member. The records are not required to be in keyed sequence.

You can specify:

```

OVRDBF      FILE(JOINAB) TOFILE(FILEA) SHARE(*YES)
OPNQRYP     FILE(FILEA FILEB) FORMAT(JOINAB) +
            JFLD((FILEA/CUST FILEB/CUST)) +
            MAPFLD((CUST 'FILEA/CUST'))
CALL        PGM(PGME) /* Created using file JOINAB as input */
CLOF        OPNID(FILEA)
DLTOVR      FILE(JOINAB)

```

File JOINAB is a physical file with no data. This is the file that contains the record format to be specified on the FORMAT parameter in the Open Query File (OPNQRYP) command.

Notice that the TOFILE parameter on the Override with Database File (OVRDBF) command specifies the name of the primary file for the join operation (the first file specified for the FILE parameter on the OPNQRYP command). In this example, the FILE parameter on the Open Query File (OPNQRYP) command identifies the files in the sequence they are to be joined (A to B). The format for the file is in the file JOINAB.

The JFLD parameter identifies the *Cust* field in FILEA to join to the *Cust* field in FILEB. Because the *Cust* field is not unique across all of the joined record formats, it must be qualified on the JFLD parameter. The system attempts to determine, in some cases, the most efficient values even if you do not specify the JFLD parameter on the Open Query File (OPNQRYP) command. For example, using the previous example, if you specified:

```

OPNQRYP     FILE(FILEA FILEB) FORMAT(JOINAB) +
            QRYSLT('FILEA/CUST *EQ FILEB/CUST') +
            MAPFLD((CUST 'FILEA/CUST'))

```

The system joins FILEA and FILEB using the *Cust* field because of the values specified for the QRYSLT parameter. Notice that in this example the JFLD parameter is not specified on the command. However, if either JDFTVAL(\*ONLYDFT) or JDFTVAL(\*YES) is specified on the OPNQRYP command, the JFLD parameter must be specified.

The MAPFLD parameter is needed on the Open Query File (OPNQRYP) command to describe which file should be used for the data for the *Cust* field in the record format for file JOINAB. If a field is defined on the MAPFLD parameter, its unqualified name (the *Cust* field in this case without the file name identification) can be used anywhere else in the OPNQRYP command. Because the *Cust* field is defined on the MAPFLD parameter, the first value of the JFLD parameter need not be qualified. For example, the same result could be achieved by specifying:

```

JFLD((CUST FILEB/CUST)) +
MAPFLD((CUST 'FILEA/CUST'))

```

Any other uses of the same field name in the Open Query File (OPNQRYP) command to indicate a field from a file other than the file defined by the MAPFLD parameter must be qualified with a file name.

Because no KEYFLD parameter is specified, the records appear in any sequence depending on how the Open Query File (OPNQRYP) command selects the records. You can force the system to arrange the records the same as the primary file. To do this, specify \*FILE on the KEYFLD parameter. You can specify this even if the primary file is in arrival sequence.



The JDFTVAL parameter (similar to the JDFTVAL keyword in DDS) can also be specified on the Open Query File (OPNQRYF) command to describe what the system should do if one of the records is missing from the secondary file. In this example, the JDFTVAL parameter was not specified, so only the records that exist in both files are selected.

If you tell the system to improve the results of the query (through parameters on the OPNQRYF command), it will generally try to use the file with the smallest number of records selected as the primary file. However, the system will also try to avoid building a temporary file.

You can force the system to follow the file sequence of the join as you have specified it in the FILE parameter on the Open Query File (OPNQRYF) command by specifying JORDER(\*FILE). If JDFTVAL(\*YES) or JDFTVAL(\*ONLYDFT) is specified, the system will never change the join file sequence because a different sequence could cause different results.

**Example 2:** Reading only those records with secondary file records

Assume you want to join files FILEAB, FILECD, and FILEEF to select only those records with matching records in secondary files. Define a file JOINF and describe the format that should be used. Assume the record formats for the files contain the following fields:

FILEAB	FILECD	FILEEF	JOINF
Abitm	Cditm	Efitm	Abitm
Abord	Cddscp	Efcolr	Abord
Abdat	Cdcolr	Efqty	Cddscp Cdcolr Efqty

In this case, all field names in the files that make up the join file begin with a 2-character prefix (identical for all fields in the file) and end with a suffix that is identical across all the files (for example, *xxitm*). This makes all field names unique and avoids having to qualify them.

The *xxitm* field allows the join from FILEAB to FILECD. The two fields *xxitm* and *xxcolr* allow the join from FILECD to FILEEF. A keyed sequence access path does not have to exist for these files. However, if a keyed sequence access path does exist, performance may improve significantly because the system will attempt to use the existing access path to arrange and select records, where it can. If access paths do not exist, the system automatically creates and maintains them as long as the file is open.

```

OVRDBF    FILE(JOINF) TOFILE(FILEAB) SHARE(*YES)
OPNQRYF   FILE(FILEAB FILECD FILEEF) +
          FORMAT(JOINF) +
          JFLD((ABITM CDITM)(CDITM EFITM) +
              (CDCOLR EFCOLR))
CALL      PGM(PGME) /* Created using file JOINF as input */
CLOF      OPNID(FILEAB)
DLTOVR    FILE(JOINF)

```

The join field pairs do not have to be specified in the order shown above. For example, the same result is achieved with a JFLD parameter value of:

```
JFLD((CDCOLR EFCOLR)(ABITM CDITM) (CDITM EFITM))
```

The attributes of each pair of join fields do not have to be identical. Normal padding of character fields and decimal alignment for numeric fields occurs automatically.

The JDFTVAL parameter is not specified so \*NO is assumed and no default values are used to construct join records. If you specified JDFTVAL(\*YES) and there is no record in file FILECD that has the same join field value as a record in file FILEAB, defaults are used for the *Cddscp* and *Cdcolr* fields to join to file FILEEF. Using these defaults, a matching record can be found in file FILEEF (depending on if the default value matches a record in the secondary file). If not, a default value appears for these files and for the *Efqty* field.

**Example 3:** Using mapped fields as join fields

You can use fields defined on the MAPFLD parameter for either one of the join field pairs. This is useful when the key in the secondary file is defined as a single field (for example, a 6-character date field) and there are separate fields for the same information (for example, month, day, and year) in the primary file. Assume FILEA has character fields *Year*, *Month*, and *Day* and needs to be joined to FILEB which has the *Date* field in YYMMDD format. Assume you have defined file JOINAB with the desired format. You can specify:

```
OVRDBF      FILE(JOINAB) TOFILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA FILEB) FORMAT(JOINAB) +
            JFLD((YYMMDD FILEB/DATE)) +
            MAPFLD((YYMMDD 'YEAR *CAT MONTH *CAT DAY'))
CALL        PGM(PGME) /* Created using file JOINAB as input */
CLOF        OPNID(FILEA)
DLTOVR      FILE(JOINAB)
```

The MAPFLD parameter defines the *YYMMDD* field as the concatenation of several fields from FILEA. You do not need to specify field attributes (for example, length or type) for the *YYMMDD* field on the MAPFLD parameter because the system calculates the attributes from the expression.

## Handling Missing Records in Secondary Join Files

The system allows you to control whether to allow defaults for missing records in secondary files (similar to the JDFTVAL DDS keyword for a join logical file). You can also specify that only records with defaults be processed. This allows you to select only those records in which there is a missing record in the secondary file.

**Example 1:** Reading records from the primary file that do not have a record in the secondary file

In Example 1 under “Dynamically Joining Database Files without DDS” on page 6-23, the JDFTVAL parameter is not specified, so the only records read are the result of a successful join from FILEA to FILEB. If you want a list of the records in FILEA that do not have a match in FILEB, you can specify \*ONLYDFT on the JDFTVAL parameter as shown in the following example:

```

OVRDBF      FILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA FILEB) FORMAT(FILEA) +
            JFLD((CUST FILEB/CUST)) +
            MAPFLD((CUST 'FILEA/CUST')) +
            JDFTVAL(*ONLYDFT)
CALL        PGM(PGME) /* Created using file FILEA as input */
CLOF        OPNID(FILEA)
DLTOVR      FILE(FILEA)

```

JDFTVAL(\*ONLYDFT) causes a record to be returned to the program only when there is no equivalent record in the secondary file (FILEB).

Because any values returned by the join operation for the fields in FILEB are defaults, it is normal to use only the format for FILEA. The records that appear are those that do not have a match in FILEB. The FORMAT parameter is required whenever the FILE parameter describes more than a single file, but the file name specified can be one of the files specified on the FILE parameter. The program is created using FILEA.

Conversely, you can also get a list of all the records where there is a record in FILEB that does not have a match in FILEA. You can do this by making the secondary file the primary file in all the specifications. You would specify:

```

OVRDBF      FILE(FILEB) SHARE(*YES)
OPNQRYF     FILE(FILEB FILEA) FORMAT(FILEB) JFLD((CUST FILEA/CUST)) +
            MAPFLD((CUST 'FILEB/CUST')) JDFTVAL(*ONLYDFT)
CALL        PGM(PGMF) /* Created using file FILEB as input */
CLOF        OPNID(FILEB)
DLTOVR      FILE(FILEB)

```

**Note:** The Override with Database File (OVRDBF) command in this example uses FILE(FILEB) because it must specify the first file on the OPNQRYF FILE parameter. The Close File (CLOF) command also names FILEB. The JFLD and MAPFLD parameters also changed. The program is created using FILEB.

## Unique-Key Processing

Unique-key processing allows you to process only the first record of a group. The group is defined by one or more records with the same set of key values. Processing the first record implies that the records you receive will have unique keys.

When you use unique-key processing, you can only read the file sequentially. The key fields are sorted according to the specified sort sequence (SRTSEQ) and language identifier (LANGID) (see “Example 3” on page 6-21 and “Example 4” on page 6-22).

### **Example 1:** Reading only unique-key records

Assume you want to process FILEA, which has records with duplicate keys for the *Cust* field. You want only the first record for each unique value of the *Cust* field to be processed by program PGMF. You can specify:

```

OVRDBF FILE(FILEA) SHARE(*YES)
OPNQRYF FILE(FILEA) KEYFLD(CUST) UNIQUEKEY(*ALL)
CALL PGM(PGMF)
CLOF OPNID(FILEA)
DLTOVR FILE(FILEA)

```

**Example 2:** Reading records using only some of the key fields

Assume you want to process the same file with the sequence: *Slsman*, *Cust*, *Date*, but you want only one record per *Slsman* and *Cust*. Assume the records in the file are:

Slsman	Cust	Date	Record #
01	5000	880109	1
01	5000	880115	2
01	4025	880103	3
01	4025	880101	4
02	3000	880101	5

You specify the number of key fields that are unique, starting with the first key field.

```

OVRDBF FILE(FILEA) SHARE(*YES)
OPNQRYF FILE(FILEA) KEYFLD(SLSMAN CUST DATE) UNIQUEKEY(2)
CALL PGM(PGMD)
CLOF OPNID(FILEA)
DLTOVR FILE(FILEA)

```

The following records are retrieved by the program:

Slsman	Cust	Date	Record #
01	4025	880101	4
01	5000	880109	1
02	3000	880101	5

**Note:** Null values are treated as equal, so only the first null value would be returned.

## Defining Fields Derived from Existing Field Definitions

Mapped field definitions:

- Allow you to create internal fields that specify selection values (see Example 7 under “Selecting Records without Using DDS” on page 6-9 for more information).
- Allow you to avoid confusion when the same field name occurs in multiple files (see Example 1 under “Dynamically Joining Database Files without DDS” on page 6-23 for more information).
- Allow you to create fields that exist only in the format to be processed, but not in the database itself. This allows you to perform translate, substring, concatenation, and complex mathematical operations. The following examples describe this function.

**Example 1:** Using derived fields

Assume you have the *Price* and *Qty* fields in the record format. You can multiply one field by the other by using the Open Query File (OPNQRYF) command to create the derived *Exten* field. You want FILEA to be processed, and you have already created FILEAA. Assume the record formats for the files contain the following fields:

FILEA	FILEAA
Order	Order
Item	Item
Qty	Exten
Price	Brfdsc
Descrp	

The *Exten* field is a mapped field. Its value is determined by multiplying *Qty* times *Price*. It is not necessary to have either the *Qty* or *Price* field in the new format, but they can exist in that format, too if you wish. The *Brfdsc* field is a brief description of the *Descrp* field (it uses the first 10 characters).

Assume you have specified PGMF to process the new format. To create this program, use FILEAA as the file to read. You can specify:

```
OVRRBF      FILE(FILEAA) TOFILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA) FORMAT(FILEAA) +
            MAPFLD((EXTEN 'PRICE * QTY') +
            (BRFDSC 'DESCRP'))
CALL        PGM(PGMF) /* Created using file FILEAA as input */
CLOF        OPNID(FILEA)
DLTOVR      FILE(FILEAA)
```

Notice that the attributes of the *Exten* field are those defined in the record format for FILEAA. If the value calculated for the field is too large, an exception is sent to the program.

It is not necessary to use the substring function to map to the *Brfdsc* field if you only want the characters from the beginning of the field. The length of the *Brfdsc* field is defined in the FILEAA record format.

All fields in the format specified on the FORMAT parameter must be described on the OPNQRYF command. That is, all fields in the output format must either exist in one of the record formats for the files specified on the FILE parameter or be defined on the MAPFLD parameter. If you have fields in the format on the FORMAT parameter that your program does not use, you can use the MAPFLD parameter to place zeros or blanks in the fields. Assume the *Fldc* field is a character field and the *Flcn* field is a numeric field in the output format, and you are using neither value in your program. You can avoid an error on the OPNQRYF command by specifying:

```
MAPFLD((FLDC ' ' ' ') (FLDN 0))
```

Notice quotation marks enclose a blank value. By using a constant for the definition of an unused field, you avoid having to create a unique format for each use of the OPNQRYF command.

### Example 2: Using built-in functions

Assume you want to calculate a mathematical function that is the sine of the *Fldm* field in FILEA. First create a file (assume it is called FILEAA) with a record format containing the following fields:

FILEA	FILEAA
Code	Code
Fldm	Fldm
	Sinm

You can then create a program (assume PGMF) using FILEAA as input and specify:

```
OVRDBF FILE(FILEAA) TOFILE(FILEA) SHARE(*YES)
OPNQRYF FILE(FILEA) FORMAT(FILEAA) +
        MAPFLD((SINM '%SIN(FLDM)'))
CALL PGM(PGMF) /* Created using file FILEAA as input */
CLOF OPNID(FILEA)
DLTOVR FILE(FILEAA)
```

The built-in function %SIN calculates the sine of the field specified as its argument. Because the *Sinm* field is defined in the format specified on the FORMAT parameter, the OPNQRYF command converts its internal definition of the sine value (in floating point) to the definition of the *Sinm* field. This technique can be used to avoid certain high-level language restrictions regarding the use of floating-point fields. For example, if you defined the *Sinm* field as a packed decimal field, PGMF could be written using any high-level language, even though the value was built using a floating-point field.

There are many other functions besides sine that can be used. Refer to the OPNQRYF command in the *CL Reference* manual for a complete list of built-in functions.

### Example 3: Using derived fields and built-in functions

Assume, in the previous example, that a field called *Fldx* also exists in FILEA, and the *Fldx* field has appropriate attributes used to hold the sine of the *Fldm* field. Also assume that you are not using the contents of the *Fldx* field. You can use the MAPFLD parameter to change the contents of a field before passing it to your high-level language program. For example, you can specify:

```
OVRDBF FILE(FILEA) SHARE(*YES)
OPNQRYF FILE(FILEA) MAPFLD((FLDX '%SIN(FLDM)'))
CALL PGM(PGMF) /* Created using file FILEA as input */
CLOF OPNID(FILEA)
DLTOVR FILE(FILEA)
```

In this case, you do not need to specify a different record format on the FORMAT parameter. (The default uses the format of the first file on the FILE parameter.) Therefore, the program is created by using FILEA. When using this technique, you must ensure that the field you redefine has attributes that allow the calculated value to process correctly. The least complicated approach is to create a separate file with the specific fields you want to process for each query.

You can also use this technique with a mapped field definition and the %XLATE function to translate a field so that it appears to the program in a different manner than what exists in the database. For example, you can translate a lowercase field so the program only sees uppercase.

| The sort sequence and language identifier can affect the results of the %MIN and  
 | %MAX built-in functions. For example, the uppercase and lowercase versions of  
 | letters can be equal or unequal depending on the selected sort sequence and lan-  
 | guage identifier. Note that the translated field value is used to determine the  
 | minimum and maximum, but the untranslated value is returned in the result record.

The example described uses FILEA as an input file. You can also update data using the OPNQRYF command. However, if you use a mapped field definition to change a field, updates to the field are ignored.

## Handling Divide by Zero

Dividing by zero is considered an error by the Open Query File (OPNQRYF) command.

Record selection is normally done before field mapping errors occur (for example, where field mapping would cause a division error). Therefore, a record can be omitted (based on the QRYSLT parameter values and valid data in the record) that would have caused a divide-by-zero error. In such an instance, the record would be omitted and processing by the OPNQRYF command would continue.

If you want a zero answer, the following describes a solution that is practical for typical commercial data.

Assume you want to divide A by B giving C (stated as  $A / B = C$ ). Assume the following definitions where B can be zero.

Field	Digits	Dec
A	6	2
B	3	0
C	6	2

The following algorithm can be used:

$(A * B) / \%MAX((B * B) .nnnn1)$

The %MAX function returns the maximum value of either  $B * B$  or a small value. The small value must have enough leading zeros so that it is less than any value calculated by  $B * B$  unless B is zero. In this example, B has zero decimal positions so .1 could be used. The number of leading zeros should be 2 times the number of decimals in B. For example, if B had 2 decimal positions, then .00001 should be used.

Specify the following MAPFLD definition:

MAPFLD((C '(A \* B) / \%MAX((B \* B) .1)'))

The intent of the first multiplication is to produce a zero dividend if B is zero. This will ensure a zero result when the division occurs. Dividing by zero does not occur if B is zero because the .1 value will be the value used as the divisor.

## Summarizing Data from Database File Records (Grouping)

The group processing function allows you to summarize data from existing database records. You can specify:

- The grouping fields
- Selection values both before and after grouping
- A keyed sequence access path over the new records
- Mapped field definitions that allow you to do such functions as sum, average, standard deviation, and variance, as well as counting the records in each group
- The sort sequence and language identifier that supply the weights by which the field values are grouped

You normally start by creating a file with a record format containing only the following types of fields:

- Grouping fields. Specified on the GRPFLD parameter that define groups. Each group contains a constant set of values for all grouping fields. The grouping fields do not need to appear in the record format identified on the FORMAT parameter.
- Aggregate fields. Defined by using the MAPFLD parameter with one or more of the following built-in functions:

%COUNT	Counts the records in a group
%SUM	A sum of the values of a field over the group
%AVG	Arithmetic average (mean) of a field, over the group
%MAX	Maximum value in the group for the field
%MIN	Minimum value in the group for the field
%STDDEV	Standard deviation of a field, over the group
%VAR	Variance of a field, over the group

- Constant fields. Allow constants to be placed in field values. The restriction that the Open Query File (OPNQRYF) command must know all fields in the output format is also true for the grouping function.

When you use group processing, you can only read the file sequentially.

### **Example 1:** Using group processing

Assume you want to group the data by customer number and analyze the amount field. Your database file is FILEA and you create a file named FILEAA containing a record format with the following fields:

FILEA	FILEAA
Cust	Cust
Type	Count (count of records per customer)
Amt	Amtsum (summation of the amount field)
	Amtavg (average of the amount field)
	Amtmax (maximum value of the amount field)

When you define the fields in the new file, you must ensure that they are large enough to hold the results. For example, if the *Amt* field is defined as 5 digits, you may want to define the *Amtsum* field as 7 digits. Any arithmetic overflow causes your program to end abnormally.



Assume the records in FILEA have the following values:

Cust	Type	Amt
001	A	500.00
001	B	700.00
004	A	100.00
002	A	1200.00
003	B	900.00
001	A	300.00
004	A	300.00
003	B	600.00

You then create a program (PGMG) using FILEAA as input to print the records.

```

OVRDBF FILE(FILEAA) TOFILE(FILEA) SHARE(*YES)
OPNQRYF FILE(FILEA) FORMAT(FILEAA) KEYFLD(CUST) +
        GRPFLD(CUST) MAPFLD((COUNT '%COUNT') +
        (AMTSUM '%SUM(AMT)') +
        (AMTAVG '%AVG(AMT)') +
        (AMTMAX '%MAX(AMT)'))
CALL PGM(PGMG) /* Created using file FILEAA as input */
CLOF OPNID(FILEA)
DLTOVR FILE(FILEAA)

```

The records retrieved by the program appear as:

Cust	Count	Amtsum	Amtavg	Amtmax
001	3	1500.00	500.00	700.00
002	1	1200.00	1200.00	1200.00
003	2	1500.00	750.00	900.00
004	2	400.00	200.00	300.00

**Note:** If you specify the GRPFLD parameter, the groups may not appear in ascending sequence. To ensure a specific sequence, you should specify the KEYFLD parameter.

Assume you want to print only the summary records in this example in which the *Amtsum* value is greater than 700.00. Because the *Amtsum* field is an aggregate field for a given customer, use the GRPSLT parameter to specify selection after grouping. Add the GRPSLT parameter:

```
GRPSLT('AMTSUM *GT 700.00')
```

The records retrieved by your program are:

Cust	Count	Amtsum	Amtavg	Amtmax
001	3	1500.00	500.00	700.00
002	1	1200.00	1200.00	1200.00
003	2	1500.00	750.00	900.00

The Open Query File (OPNQRYF) command supports selection both before grouping (QRYSLT parameter) and after grouping (GRPSLT parameter).

Assume you want to select additional customer records in which the *Type* field is equal to A. Because *Type* is a field in the record format for file FILEA and not an aggregate field, you add the QRYSLT statement to select before grouping as follows:

```
QRYSLT('TYPE *EQ "A" ')
```

Note that fields used for selection do not have to appear in the format processed by the program.

The records retrieved by your program are:

Cust	Count	Amtsum	Amtavg	Amtmax
001	2	800.00	400.00	500.00
002	1	1200.00	1200.00	1200.00

Notice the values for CUST 001 changed because the selection took place before the grouping took place.

Assume you want to arrange the output by the *Amtavg* field in descending sequence, in addition to the previous QRYSLT parameter value. You can do this by changing the KEYFLD parameter on the OPNQRYF command as:

```
KEYFLD((AMTAVG *DESCEND))
```

The records retrieved by your program are:

Cust	Count	Amtsum	Amtavg	Amtmax
002	1	1200.00	1200.00	1200.00
001	2	800.00	400.00	500.00

## Final Total-Only Processing

Final-total-only processing is a special form of grouping in which you do not specify grouping fields. Only one record is output. All of the special built-in functions for grouping can be specified. You can also specify the selection of records that make up the final total.

### **Example 1:** Simple total processing

Assume you have a database file FILEA and decide to create file FINTOT for your final total record as follows:

FILEA	FINTOT
Code	Count (count of all the selected records)
Amt	Totamt (total of the amount field)
	Maxamt (maximum value in the amount field)

The FINTOT file is created specifically to hold the single record which is created with the final totals. You would specify:

```

OVRDBF      FILE(FINTOT) TOFILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA) FORMAT(FINTOT) +
            MAPFLD((COUNT '%COUNT') +
            (TOTAMT '%SUM(AMT)') (MAXAMT '%MAX(AMT)'))
CALL        PGM(PGMG) /* Created using file FINTOT as input */
CLOF        OPNID(FILEA)
DLTOVR      FILE(FINTOT)

```

**Example 2:** Total-only processing with record selection

Assume you want to change the previous example so that only the records where the *Code* field is equal to B are in the final total. You can add the QRYSLT parameter as follows:

```

OVRDBF      FILE(FINTOT) TOFILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA) FORMAT(FINTOT) +
            QRYSLT('CODE *EQ "B" ') MAPFLD((COUNT '%COUNT') +
            (TOTAMT '%SUM(AMT)') (MAXAMT '%MAX(AMT)'))
CALL        PGM(PGMG) /* Created using file FINTOT as input */
CLOF        OPNID(FILEA)
DLTOVR      FILE(FINTOT)

```

You can use the GRPSLT keyword with the final total function. The GRPSLT selection values you specify determines if you receive the final total record.

**Example 3:** Total-only processing using a new record format

Assume you want to process the new file/format with a CL program. You want to read the file and send a message with the final totals. You can specify:

```

DCLF        FILE(FINTOT)
DCL         &COUNTA *CHAR LEN(7)
DCL         &TOTAMTA *CHAR LEN(9)
OVRDBF      FILE(FINTOT) TOFILE(FILEA) SHARE(*YES)
OPNQRYF     FILE(FILEA) FORMAT(FINTOT) MAPFLD((COUNT '%COUNT') +
            (TOTAMT '%SUM(AMT)'))
RCVF
CLOF        OPNID(FILEA)
CHGVAR      &COUNTA &COUNT
CHGVAR      &TOTAMTA &TOTAMT
SNDPGMMSG   MSG('COUNT=' *CAT &COUNTA *CAT +
            ' Total amount=' *CAT &TOTAMTA)
DLTOVR      FILE(FINTOT)

```

You must convert the numeric fields to character fields to include them in an immediate message.

## Controlling How the System Runs the Open Query File Command

The optimization function allows you to specify how you are going to use the results of the query.

When you use the Open Query File (OPNQRYF) command there are two steps where performance considerations exist. The first step is during the actual processing of the OPNQRYF command itself. This step decides if OPNQRYF is going to use an existing access path or build a new one for this query request. The second step when performance considerations play a role is when the application program is using the results of the OPNQRYF to process the data. (See

Appendix D, "Design Guidelines for OPNQRYF Performance" for additional design guidelines.)

For most batch type functions, you are usually only interested in the total time of both steps mentioned above. Therefore, the default for OPNQRYF is OPTIMIZE(\*ALLIO). This means that OPNQRYF will consider the total time it takes for both steps.

If you use OPNQRYF in an interactive environment, you may not be interested in processing the entire file. You may want the first screen full of records to be displayed as quickly as possible. For this reason, you would want the first step to avoid building an access path, if possible. You might specify OPTIMIZE(\*FIRSTIO) in such a situation.

If you want to process the same results of OPNQRYF with multiple programs, you would want the first step to make an efficient open data path (ODP). That is, you would try to minimize the number of records that must be read by the processing program in the second step by specifying OPTIMIZE(\*MINWAIT) on the OPNQRYF command.

If the KEYFLD or GRPFLD parameters on the OPNQRYF command require that an access path be built when there is no access path to share, the access path is built entirely regardless of the OPTIMIZE entry. Optimization mainly affects selection processing.

**Example 1:** Optimizing for the first set of records

Assume that you have an interactive job in which the operator requests all records where the *Code* field is equal to B. Your program's subfile contains 15 records per screen. You want to get the first screen of results to the operator as quickly as possible. You can specify:

```
OVRDBF FILE(FILEA) SHARE(*YES)
OPNQRYF FILE(FILEA) QRYSLT('CODE = "B" ') +
        SEQONLY(*YES 15) OPTIMIZE(*FIRSTIO)
CALL PGM(PGMA)
CLOF OPNID(FILEA)
DLTOVR FILE(FILEA)
```

The system optimizes handling the query and fills the first buffer with records before completing the entire query regardless of whether an access path already exists over the *Code* field.

**Example 2:** Optimizing to minimize the number of records read

Assume that you have multiple programs that will access the same file which is built by the Open Query File (OPNQRYF) command. In this case, you will want to optimize the performance so that the application programs read only the data they are interested in. This means that you want OPNQRYF to perform the selection as efficiently as possible. You could specify:

```

OVRDBF  FILE(FILEA) SHARE(*YES)
OPNQRYF FILE(FILEA) QRYSLT('CODE *EQ "B" ') +
        KEYFLD(CUST) OPTIMIZE(*MINWAIT)
CALL    PGM(PGMA)
POSDBF  OPNID(FILEA) POSITION(*START)
CALL    PGM(PGMB)
CLOF    OPNID(FILEA)
DLTOVR  FILE(FILEA)

```

## Considerations for Creating a File and Using the FORMAT Parameter

You must specify a record format name on the FORMAT parameter when you request join processing by specifying multiple entries on the FILE parameter (that is, you cannot specify FORMAT(\*FILE)). Also, a record format name is normally specified with the grouping function or when you specify a complex expression on the MAPFLD parameter to define a derived field. Consider the following:

- The record format name is any name you select. It can differ from the format name in the database file you want to query.
- The field names are any names you select. If the field names are unique in the database files you are querying, the system implicitly maps the values for any fields with the same name in a queried file record format (FILE parameter) and in the query result format (FORMAT parameter). See Example 1 under “Dynamically Joining Database Files without DDS” on page 6-23 for more information.
- If the field names are unique, but the attributes differ between the file specified on the FILE parameter and the file specified on the FORMAT parameter, the data is implicitly mapped.
- The correct field attributes must be used when using the MAPFLD parameter to define derived fields. For example, if you are using the grouping %SUM function, you must define a field that is large enough to contain the total. If not, an arithmetic overflow occurs and an exception is sent to the program.
- Decimal alignment occurs for all field values mapped to the record format identified on the FORMAT parameter. Assume you have a field in the query result record format with 5 digits with 0 decimals, and the value that was calculated or must be mapped to that field is 0.12345. You will receive a result of 0 in your field because digits to the right of the decimal point are truncated.

## Considerations for Arranging Records

The default processing for the Open Query File (OPNQRYF) command provides records in any order that improves performance and does not conflict with the order specified on the KEYFLD parameter. Therefore, unless you specify the KEYFLD parameter to either name specific key fields or specify KEYFLD(\*FILE), the sequence of the records returned to your program can vary each time you run the same Open Query File (OPNQRYF) command.

When you specify the KEYFLD(\*FILE) parameter option for the Open Query File (OPNQRYF) command, and a sort sequence other than \*HEX has been specified for the query with the job default or the OPNQRYF SRTSEQ parameter, you can receive your records in an order that does not reflect the true file order. If the file is keyed, the query's sort sequence is applied to the key fields of the file and informational message CPI431F is sent. The file's sort sequence and alternative collating sequence table are ignored for the ordering, if they exist. This allows users to indi-

cate which fields to apply a sort sequence to without having to list all the field names. If a sort sequence is not specified for the query (for example, \*HEX), ordering is done as it was prior to Version 2 Release 3.

## Considerations for DDM Files

The Open Query File (OPNQRYF) command can process DDM files. All files identified on the FILE parameter must exist on the same IBM AS/400 system or System/38 target system. An OPNQRYF which specifies group processing and uses a DDM file requires that both the source and target system be the same type (either both System/38 or both AS/400 systems).

## Considerations for Writing a High-Level Language Program

For the method described under “Using an Existing Record Format in the File” on page 6-4 (where the FORMAT parameter is omitted), your high-level language program is coded as if you are directly accessing the database file. Selection or sequencing occurs external to your program, and the program receives the selected records in the order you specified. The program does not receive records that are omitted by your selection values. This same function occurs if you process through a logical file with select/omit values.

If you use the FORMAT parameter, your program specifies the same file name used on the FORMAT parameter. The program is written as if this file contains actual data.

If you read the file sequentially, your high-level language can automatically specify that the key fields are ignored. Normally you write the program as if it is reading records in arrival sequence. If the KEYFLD parameter is used on the Open Query File (OPNQRYF) command, you receive a diagnostic message, which can be ignored.

If you process the file randomly by keys, your high-level language probably requires a key specification. If you have selection values, it can prevent your program from accessing a record that exists in the database. A Record not found condition can occur on a random read whether the OPNQRYF command was used or whether a logical file created using DDS select/omit logic was used.

In some cases, you can monitor exceptions caused by mapping errors such as arithmetic overflow, but it is better to define the attributes of all fields to correctly handle the results.

## Messages Sent When the Open Query File (OPNQRYF) Command Is Run

When the OPNQRYF command is run, messages are sent informing the interactive user of the status of the OPNQRYF request. For example, a message would be sent to the user if a keyed access path was built by the OPNQRYF to satisfy the request. The following messages might be sent during a run of the OPNQRYF command:

Message Identifier	Description
CPI4301	Query running.

Message Identifier	Description
CPI4302	Query running. Building access path...
CPI4303	Query running. Creating copy of file...
CPI4304	Query running. Selection complete...
CPI4305	Query running. Sorting copy of file...
CPI4306	Query running. Building access path from file...
CPI4011	Query running. Number of records processed...

To stop these status messages from appearing, see the discussion about message handling in the *CL Programmer's Guide*.

When your job is running under debug (using the STRDBG command), messages are sent to your job log that describe the implementation method used to process the OPNQRYF request. These messages provide information about the optimization processing that occurred. They can be used as a tool for tuning the OPNQRYF request to achieve the best performance. The messages are as follows:

CPI4321 Access path built for file...  
 CPI4322 Access path built from keyed file...  
 CPI4324 Temporary file built from file...  
 CPI4325 Temporary file built for query  
 CPI4326 File processed in join position...  
 CPI4327 File processed in join position 1.  
 CPI4328 Access path of file used...  
 CPI4329 Arrival sequence used for file...  
 CPI432A Query optimizer timed out...  
 CPI432C All access paths were considered for file...

Most of the messages provide a reason why the particular option was performed. The second level text on each message gives an extended description of why the option was chosen. Some messages provide suggestions to help improve the performance of the OPNQRYF request.

## Using the Open Query File (OPNQRYF) Command for More Than Just Input

The OPNQRYF command supports the OPTION parameter to determine the type of processing. The default is OPTION(\*INP), so the file is opened for input only. You can also use other OPTION values on the OPNQRYF command and a high-level language program to add, update, or delete records through the open query file. However, if you specify the UNIQUEKEY, GRPFLD, or GRPSLT parameters, use one of the aggregate functions, or specify multiple files on the FILE parameter, your use of the file is restricted to input only.

A join logical file is limited to input-only processing. A view is limited to input-only processing, if group, join, union, or distinct processing is specified in the definition of the view.

If you want to change a field value from the current value to a different value in some of the records in a file, you can use a combination of the OPNQRYF

command and a specific high-level language program. For example, assume you want to change all the records where the *Flda* field is equal to ABC so that the *Flda* field is equal to XYZ. You can specify:

```
OVRDBF FILE(FILEA) SHARE(*YES)
OPNQRYF FILE(FILEA) OPTION(*ALL) QRYSLT('FLDA *EQ "ABC" ')
CALL PGM(PGMA)
CLOF OPNID(FILEA)
DLTOVR FILE(FILEA)
```

Program PGMA processes all records it can read, but the query selection restricts these to records where the *Flda* field is equal to ABC. The program changes the field value in each record to XYZ and updates the record.

You can also delete records in a database file using the OPNQRYF command. For example, assume you have a field in your record that, if equal to X, means the record should be deleted. Your program can be written to delete any records it reads and use the OPNQRYF command to select those to be deleted such as:

```
OVRDBF FILE(FILEA) SHARE(*YES)
OPNQRYF FILE(FILEA) OPTION(*ALL) QRYSLT('DLTCOD *EQ "X" ')
CALL PGM(PGMB)
CLOF OPNID(FILEA)
DLTOVR FILE(FILEA)
```

You can also add records by using the OPNQRYF command. However, if the query specifications include selection values, your program can be prevented from reading the added records because of the selection values.

## Date, Time, and Timestamp Comparisons Using the OPNQRYF Command

A date, time, or timestamp value can be compared either with another value of the same data type or with a string representation of that data type. All comparisons are chronological, which means the farther a time is from January 1, 0001, the *greater* the value of that time.

Comparisons involving time values and string representations of time values always include seconds. If the string representation omits seconds, zero seconds are implied.

Comparisons involving timestamp values are chronological without regard to representations that might be considered equivalent. Thus, the following predicate is true:

```
TIMESTAMP('1990-02-23-00.00.00') > '1990-02-22-24.00.00'
```

When a character, DBCS-open, or DBCS-either field or constant is represented as a date, time, or timestamp, the following rules apply:

**Date:** The length of the field or literal must be at least 8 if the date format is \*ISO, \*USA, \*EUR, \*JIS, \*YMD, \*MDY, or \*DMY. If the date format is \*JUL (yyddd), the length of the variable must be at least 6 (includes the separator between yy and ddd). The field or literal may be padded with blanks.

**Time:** For all of the time formats (\*USA, \*ISO, \*EUR, \*JIS, \*HMS), the length of the field or literal must be at least 4. The field or literal may be padded with blanks.



**Timestamp:** For the timestamp format (yyyy-mm-dd-hh.mm.ss.uuuuuu), the length of the field or literal must be at least 16. The field or literal may be padded with blanks.

## Date, Time, and Timestamp Arithmetic Using OPNQRYP CL Command

Date, time, and timestamp values can be incremented, decremented, and subtracted. These operations may involve decimal numbers called *durations*. Following is a definition of durations and a specification of the rules for performing arithmetic operations on date, time, and timestamp values.

### Durations

A **duration** is a number representing an interval of time. The four types of durations are:

#### Labeled Duration

A **labeled duration** represents a specific unit of time as expressed by a number (which can be the result of an expression) used as an operand for one of the seven duration built-in functions: %DURYEAR, %DURMONTH, %DURDAY, %DURHOUR, %DURMINUTE, %DURSEC, or %DURMICSEC. The functions are for the duration of year, month, day, hour, minute, second, and microsecond, respectively. The number specified is converted as if it was assigned to a DECIMAL(15,0) number. A labeled duration can only be used as an operand of an arithmetic operator when the other operand is a value of data type \*DATE, \*TIME, or \*TIMESTP. Thus, the expression HIREDATE + %DURMONTH(2) + %DURDAY(14) is valid, whereas the expression HIREDATE + (%DURMONTH(2) + %DURDAY(14)) is not. In both of these expressions, the labeled durations are %DURMONTH(2) and %DURDAY(14).

#### Date Duration

A **date duration** represents a number of years, months, and days, expressed as a DECIMAL(8,0) number. To be properly interpreted, the number must have the format *yyyymmdd*, where *yyyy* represents the number of years, *mm* the number of months, and *dd* the number of days. The result of subtracting one date value from another, as in the expression HIREDATE - BRTHDATE, is a date duration.

#### Time Duration

A **time duration** represents a number of hours, minutes, and seconds, expressed as a DECIMAL(6,0) number. To be properly interpreted, the number must have the format *hhmmss*, where *hh* represents the number of hours, *mm* the number of minutes, and *ss* the number of seconds. The result of subtracting one time value from another is a time duration.

#### Timestamp Duration

A **timestamp duration** represents a number of years, months, days, hours, minutes, seconds, and microseconds, expressed as a DECIMAL(20,6) number. To be properly interpreted, the number must have the format *yyyymmddhhmmsszzzzzz*, where *yyyy*, *mm*, *dd*, *hh*, *mm*, *ss*, and *zzzzzz* represent, respectively, the number of years, months, days, hours, minutes, seconds, and microseconds. The result of subtracting one timestamp value from another is a timestamp duration.

## Rules for Date, Time, and Timestamp Arithmetic

The only arithmetic operations that can be performed on date and time values are addition and subtraction. If a date or time value is the operand of addition, the other operand must be a duration. The specific rules governing the use of the addition operator with date and time values follow:

- If one operand is a date, the other operand must be a date duration or a labeled duration of years, months, or days.
- If one operand is a time, the other operand must be a time duration or a labeled duration of hours, minutes, or seconds.
- If one operand is a timestamp, the other operand must be a duration. Any type of duration is valid.

The rules for the use of the subtraction operator on date and time values are not the same as those for addition because a date or time value cannot be subtracted from a duration, and because the operation of subtracting two date and time values is not the same as the operation of subtracting a duration from a date or time value. The specific rules governing the use of the subtraction operator with date and time values follow:

- If the first operand is a date, the second operand must be a date, a date duration, a string representation of a date, or a labeled duration of years, months, or days.
- If the second operand is a date, the first operand must be a date or a string representation of a date.
- If the first operand is a time, the second operand must be a time, a time duration, a string representation of a time, or a labeled duration of hours, minutes, or seconds.
- If the second operand is a time, the first operand must be a time or string representation of a time.
- If the first operand is a timestamp, the second operand must be a timestamp, a string representation of a timestamp, or a duration.
- If the second operand is a timestamp, the first operand must be a timestamp or a string representation of a timestamp.

## Date Arithmetic

Dates can be subtracted, incremented, or decremented.

**Subtracting Dates:** The result of subtracting one date (DATE2) from another (DATE1) is a date duration that specifies the number of years, months, and days between the two dates. The data type of the result is DECIMAL(8,0). If DATE1 is greater than or equal to DATE2, DATE2 is subtracted from DATE1. If DATE1 is less than DATE2, however, DATE1 is subtracted from DATE2, and the sign of the result is made negative. The following procedural description clarifies the steps involved in the operation  $RESULT = DATE1 - DATE2$ .

```
If %DAY(DATE2) <= %DAY(DATE1)
;   then %DAY(RESULT) = %DAY(DATE1) - %DAY(DATE2).

If %DAY(DATE2) > %DAY(DATE1)
;   then %DAY(RESULT) = N + %DAY(DATE1) - %DAY(DATE2)
;   where N = the last day of %MONTH(DATE2).
;   %MONTH(DATE2) is then incremented by 1.
```

```

If %MONTH(DATE2) <= %MONTH(DATE1)
;   then %MONTH(RESULT) = %MONTH(DATE1) - %MONTH(DATE2).
If %MONTH(DATE2) > %MONTH(DATE1)
;   then %MONTH(RESULT) = 12 + %MONTH(DATE1) - %MONTH(DATE2).
;   %YEAR(RESULT) is then incremented by 1.
%YEAR(RESULT) = %YEAR(DATE1) - %YEAR(DATE2).

```

For example, the result of %DATE('3/15/2000') - '12/31/1999' is 215 (or, a duration of 0 years, 2 months, and 15 days).

**Incrementing and Decrementing Dates:** The result of adding a duration to a date, or of subtracting a duration from a date, is itself a date. (For the purposes of this operation, a month denotes the equivalent of a calendar page. Adding months to a date, then, is like turning the pages of a calendar, starting with the page on which the date appears.) The result must fall between the dates January 1, 0001, and December 31, 9999, inclusive. If a duration of years is added or subtracted, only the year portion of the date is affected. The month is unchanged, as is the day unless the result would be February 29 of a year that is not a leap year. In this case, the day is changed to 28.

Similarly, if a duration of months is added or subtracted, only months and, if necessary, years are affected. The day portion of the date is unchanged unless the result would not be valid (September 31, for example). In this case, the day is set to the last day of the month.

Adding or subtracting a duration of days will, of course, affect the day portion of the date, and potentially the month and year.

Date durations, whether positive or negative, may also be added to and subtracted from dates. As with labeled durations, the result is a valid date.

When a positive date duration is added to a date, or a negative date duration is subtracted from a date, the date is incremented by the specified number of years, months, and days, in that order. Thus, DATE1 + X, where X is a positive DECIMAL(8,0) number, is equivalent to the expression: DATE1 + %DURYEAR(%YEAR(X)) + %DURMONTH(%MONTH(X)) + %DURDAY(%DAY(X))

When a positive date duration is subtracted from a date, or a negative date duration is added to a date, the date is decremented by the specified number of days, months, and years, in that order. Thus, DATE1 - X, where X is a positive DECIMAL(8,0) number, is equivalent to the expression: DATE1 - %DURDAY(%DAY(X)) - %DURMONTH(%MONTH(X)) - %DURYEAR(%YEAR(X))

When adding durations to dates, adding one month to a given date gives the same date one month later *unless* that date does not exist in the later month. In that case, the date is set to that of the last day of the later month. For example, January 28 plus one month gives February 28; and one month added to January 29, 30, or 31 results in either February 28 or, for a leap year, February 29.

**Note:** If one or more months are added to a given date and then the same number of months is subtracted from the result, the final date is not necessarily the same as the original date.

## Time Arithmetic

Times can be subtracted, incremented, or decremented.

**Subtracting Times:** The result of subtracting one time (TIME2) from another (TIME1) is a time duration that specifies the number of hours, minutes, and seconds between the two times. The data type of the result is DECIMAL(6,0). If TIME1 is greater than or equal to TIME2, TIME2 is subtracted from TIME1. If TIME1 is less than TIME2, however, TIME1 is subtracted from TIME2, and the sign of the result is made negative. The following procedural description clarifies the steps involved in the operation  $RESULT = TIME1 - TIME2$ .

```
If %SECOND(TIME2) <= %SECOND(TIME1)
;   then %SECOND(RESULT) = %SECOND(TIME1) - %SECOND(TIME2).

If %SECOND(TIME2) > %SECOND(TIME1)
;   then %SECOND(RESULT) = 60 + %SECOND(TIME1) -
%SECOND(TIME2).
;   %MINUTE(TIME2) is then incremented by 1.

If %MINUTE(TIME2) <= %MINUTE(TIME1)
;   then %MINUTE(RESULT) = %MINUTE(TIME1) - %MINUTE(TIME2).

If %MINUTE(TIME2) > %MINUTE(TIME1)
;   then %MINUTE(RESULT) = 60 + %MINUTE(TIME1) - %MINUTE(TIME2).
;   %HOUR(TIME2) is then incremented by 1.

%HOUR(RESULT) = %HOUR(TIME1) - %HOUR(TIME2).
```

For example, the result of  $\%TIME('11:02:26') - '00:32:56'$  is 102930 (a duration of 10 hours, 29 minutes, and 30 seconds).

**Incrementing and Decrementing Times:** The result of adding a duration to a time, or of subtracting a duration from a time, is itself a time. Any overflow or underflow of hours is discarded, thereby ensuring that the result is always a time. If a duration of hours is added or subtracted, only the hours portion of the time is affected. The minutes and seconds are unchanged.

Similarly, if a duration of minutes is added or subtracted, only minutes and, if necessary, hours are affected. The seconds portion of the time is unchanged.

Adding or subtracting a duration of seconds will, of course, affect the seconds portion of the time, and potentially the minutes and hours.

Time durations, whether positive or negative, also can be added to and subtracted from times. The result is a time that has been incremented or decremented by the specified number of hours, minutes, and seconds, in that order.  $TIME1 + X$ , where  $X$  is a DECIMAL(6,0) number, is equivalent to the expression:  $TIME1 + \%DURHOUR(\%HOUR(X)) + \%DURMINUTE(\%MINUTE(X)) + \%DURSEC(\%SECOND(X))$

## Timestamp Arithmetic

Timestamps can be subtracted, incremented, or decremented.

**Subtracting Timestamps:** The result of subtracting one timestamp (TS2) from another (TS1) is a timestamp duration that specifies the number of years, months, days, hours, minutes, seconds, and microseconds between the two timestamps. The data type of the result is DECIMAL(20,6). If TS1 is greater than or equal to

TS2, TS2 is subtracted from TS1. If TS1 is less than TS2, however, TS1 is subtracted from TS2 and the sign of the result is made negative. The following procedural description clarifies the steps involved in the operation  $RESULT = TS1 - TS2$ :

```
If %MICSEC(TS2) <= %MICSEC(TS1)
;   then %MICSEC(RESULT) = %MICSEC(TS1) -
;   %MICSEC(TS2).

If %MICSEC(TS2) > %MICSEC(TS1)
;   then %MICSEC(RESULT) = 1000000 +
;   %MICSEC(TS1) - %MICSEC(TS2)
;   and %SECOND(TS2) is incremented by 1.
```

The seconds and minutes part of the timestamps are subtracted as specified in the rules for subtracting times:

```
If %HOUR(TS2) <= %HOUR(TS1)
;   then %HOUR(RESULT) = %HOUR(TS1) - %HOUR(TS2).

If %HOUR(TS2) > %HOUR(TS1)
;   then %HOUR(RESULT) = 24 + %HOUR(TS1) - %HOUR(TS2)
;   and %DAY(TS2) is incremented by 1.
```

The date part of the timestamp is subtracted as specified in the rules for subtracting dates.

**Incrementing and Decrementing Timestamps:** The result of adding a duration to a timestamp, or of subtracting a duration from a timestamp, is itself a timestamp. Date and time arithmetic is performed as previously defined, except that an overflow or underflow of hours is carried into the date part of the result, which must be within the range of valid dates. Microseconds overflow into seconds.

## Using the Open Query File (OPNQRYF) Command for Random Processing

Most of the previous examples show the OPNQRYF command using sequential processing. Random processing operations (for example, the RPG/400 language operation CHAIN or the COBOL/400 language operation READ) can be used in most cases. However, if you are using the group or unique-key functions, you cannot process the file randomly.

## Performance Considerations

See Appendix D, "Design Guidelines for OPNQRYF Performance" for design guidelines, tips, and techniques for optimizing the performance of a query application.

The best performance can occur when the Open Query File (OPNQRYF) command uses an existing keyed sequence access path. For example, if you want to select all the records where the *Code* field is equal to B and an access path exists over the *Code* field, the system can use the access path to perform the selection (key positioning selection) rather than read the records and select at run time (dynamic selection).

The Open Query File (OPNQRYF) command cannot use an existing index when any of the following are true:

- The key field in the access path is derived from a substring function.
- The key field in the access path is derived from a concatenation function.

- Both of the following are true of the sort sequence table associated with the query (specified on the SRTSEQ parameter):
  - It is a shared-weight sequence table.
  - It does not match the sequence table associated with the access path (a sort sequence table or an alternate collating sequence table).
- Both of the following are true of the sort sequence table associated with the query (specified on the SRTSEQ parameter):
  - It is a unique-weight sequence table.
  - It does not match the sequence table associated with the access path (a sort sequence table or an alternate collating sequence table) when either:
    - Ordering is specified (KEYFLD parameter).
    - Record selection exists (QRYSLT parameter) that does not use \*EQ, \*NE, \*CT, %WLDCRD, or %VALUES.
    - Join selection exists (JFLD parameter) that does not use \*EQ or \*NE operators.

Part of the OPNQRYF processing is to determine what is the fastest approach to satisfying your request. If the file you are using is large and most of the records have the *Code* field equal to B, it is faster to use arrival sequence processing than to use an existing keyed sequence access path. Your program will still see the same records. OPNQRYF can only make this type of decision if an access path exists on the *Code* field. In general, if your request will include approximately 20% or more of the number of records in the file, OPNQRYF will tend to ignore the existing access paths and read the file in arrival sequence.

If no access path exists over the *Code* field, the program reads all of the records in the file and passes only the selected records to your program. That is, the file is processed in arrival sequence.

The system can perform selection faster than your application program. If no appropriate keyed sequence access path exists, either your program or the system makes the selection of the records you want to process. Allowing the system to perform the selection process is considerably faster than passing all the records to your application program.

This is especially true if you are opening a file for update operations because individual records must be passed to your program, and locks are placed on every record read (in case your program needs to update the record). By letting the system perform the record selection, the only records passed to your program and locked are those that meet your selection values.

If you use the KEYFLD parameter to request a specific sequence for reading records, the fastest performance results if an access path already exists that uses the same key specification or if a keyed sequence access path exists that is similar to your specifications (such as a key that contains all the fields you specified plus some additional fields on the end of the key). This is also true for the GRPFLD parameter and on the to-fields of the JFLD parameter. If no such access path exists, the system builds an access path and maintains it as long as the file is open in your job.

Processing all of the records in a file by an access path that does not already exist is generally not as efficient as using a full record sort, if the number of records to be arranged (not necessarily the total number of records in the file) exceeds 1000.

While it is generally faster to build the keyed sequence access path than to do the sort, faster processing allowed by the use of arrival sequence processing normally favors sorting the data when looking at the total job time. If a usable access path already exists, using the access path can be faster than sorting the data. You can use the ALWCPYDTA(\*OPTIMIZE) parameter of the Open Query File (OPNQRYF) command to allow the system to use a full record sort if that is the fastest method of processing records.

If you do not intend to read all of the query records and if the OPTIMIZE parameter is \*FIRSTIO or \*MINWAIT, you can specify a number to indicate how many records you intend to retrieve. If the number of records is considerably less than the total number the query is expected to return, the system may select a faster access method.

If you use the grouping function, faster performance is achieved if you specify selection before grouping (QRYSLT parameter) instead of selection after grouping (GRPSLT parameter). Only use the GRPSLT parameter for comparisons involving aggregate functions.

For most uses of the OPNQRYF command, new or existing access paths are used to access the data and present it to your program. In some cases of the OPNQRYF command, the system must create a temporary file. The rules for when a temporary file is created are complex, but the following are typical cases in which this occurs:

- When you specify a dynamic join, and the KEYFLD parameter describes key fields from different physical files.
- When you specify a dynamic join and the GRPFLD parameter describes fields from different physical files.
- When you specify both the GRPFLD and KEYFLD parameters but they are not the same.
- When the fields specified on the KEYFLD parameter total more than 2000 bytes in length.
- When you specify a dynamic join and \*MINWAIT for the OPTIMIZE parameter.
- When you specify a dynamic join using a join logical file and the join type (JDFTVAL) of the join logical file does not match the join type of the dynamic join.
- When you specify a logical file and the format for the logical file refers to more than one physical file.
- When you specify an SQL view, the system may require a temporary file to contain the results of the view.
- When the ALWCPYDTA(\*OPTIMIZE) parameter is specified and using a temporary result would improve the performance of the query.

When a dynamic join occurs (JDFTVAL(\*NO)), OPNQRYF attempts to improve performance by reordering the files and joining the file with the smallest number of selected records to the file with the largest number of selected records. To prevent OPNQRYF from reordering the files, specify JORDER(\*FILE). This forces OPNQRYF to join the files in the sequence specify in the OPNQRYF command.

## Performance Considerations for Sort Sequence Tables

### Grouping, Joining, and Selection

When using an existing index, the optimizer ensures that the attributes of the selection, join, and grouping fields match the attributes of the keys in the existing index. Also, the sort sequence table associated with the query must match the sequence table (a sort sequence table or an alternate collating sequence table) associated with the key field of the existing index. If the sequence tables do not match, the existing index cannot be used.

However, if the sort sequence table associated with the query is a unique-weight sequence table (including \*HEX), some additional optimization is possible. The optimizer acts as though no sort sequence table is specified for any grouping fields or any selection or join predicates that use the following operators or functions:

- \*EQ
- \*NE
- \*CT
- %WLDCRD
- %VALUES

The advantage is that the optimizer is free to use any existing access path where the keys match the field and the access path either:

- Does not contain a sequence table.
- Contains a unique-weight sequence table (the table does not have to match the unique-weight sort sequence table associated with the query).

### Ordering

For ordering fields, the optimizer is not free to use any existing access path. The sort sequence tables associated with the index and the query must match unless the optimizer chooses to do a sort to satisfy the ordering request. When a sort is used, the translation is performed during the sort, leaving the optimizer free to use any existing access path that meets the selection criteria.

### Examples

In the following examples, assume that three access paths (indices) exist over the JOB field. These access paths use the following sort sequence tables:

1. SRTSEQ(\*HEX)
2. SRTSEQ(\*LANGIDUNQ) LANGID(ENU)
3. SRTSEQ(\*LANGIDSHR) LANGID(ENU)

**Example 1:** EQ selection with no sequence table.

```
OPNQRYF FILE(STAFF) QRYSLT('JOB *EQ "MGR"') SRTSEQ(*HEX)
```

The optimizer can use index 1 (\*HEX) or 2 (\*LANGIDUNQ).

**Example 2:** EQ selection with unique-weight sequence table.

```
OPNQRYF FILE(STAFF) QRYSLT('JOB *EQ "MGR"') SRTSEQ(*LANGIDUNQ) LANGID(ENU)
```

The optimizer can use index 1 (\*HEX) or 2 (\*LANGIDUNQ).

**Example 3:** EQ selection with shared-weight sequence table.



| OPNQRYF FILE(STAFF) QRYSLT('JOB \*EQ "MGR"') SRTSEQ(\*LANGIDSHR) LANGID(ENU)

| The optimizer can only use index 3 (\*LANGIDSHR).

| **Example 4:** GT selection with unique-weight sequence table.

| OPNQRYF FILE(STAFF) QRYSLT('JOB \*GT "MGR"') SRTSEQ(\*LANGIDUNQ) LANGID(ENU)

| The optimizer can only use index 2 (\*LANGIDUNQ).

| **Example 5:** Join selection with unique-weight sequence table.

| OPNQRYF FILE((STAFF1)(STAFF2)) JFLD(1/JOB 2/JOB \*EQ)  
| SRTSEQ(\*LANGIDUNQ) LANGID(ENU)

| The optimizer can use index 1 (\*HEX) or 2 (\*LANGIDUNQ).

| **Example 6:** Join selection with shared-weight sequence table.

| OPNQRYF FILE((STAFF1)(STAFF2)) JFLD(1/JOB 2/JOB \*EQ)  
| SRTSEQ(\*LANGIDSHR) LANGID(ENU)

| The optimizer can only use index 3 (\*LANGIDSHR).

| **Example 7:** Ordering with no sequence table.

| OPNQRYF FILE(STAFF) QRYSLT('JOB \*EQ "MGR"')  
| SRTSEQ(\*HEX) KEYFLD(JOB)

| The optimizer can only use index 1 (\*HEX).

| **Example 8:** Ordering with unique-weight sequence table.

| OPNQRYF FILE(STAFF) QRYSLT('JOB \*EQ "MGR"')  
| SRTSEQ(\*LANGIDUNQ) LANGID(ENU) KEYFLD(JOB)

| The optimizer can only use index 2 (\*LANGIDUNQ).

| **Example 9:** Ordering with shared-weight sequence table.

| OPNQRYF FILE(STAFF) QRYSLT('JOB \*EQ "MGR"')  
| SRTSEQ(\*LANGIDSHR) LANGID(ENU) KEYFLD(JOB)

| The optimizer can only use index 3 (\*LANGIDSHR).

| **Example 10:** Ordering with ALWCPYDTA and unique-weight sequence table.

| OPNQRYF FILE(STAFF) QRYSLT('JOB \*EQ "MGR"')  
| SRTSEQ(\*LANGIDUNQ) LANGID(ENU) KEYFLD(JOB)  
| ALWCPYDTA(\*OPTIMIZE)

| The optimizer can use index 1 (\*HEX) or 2 (\*LANGIDUNQ) for selection. Ordering  
| is done during the sort using a \*LANGIDUNQ sequence table.

| **Example 11:** Grouping with no sequence table

| OPNQRYF FILE(STAFF) GRPFLD(JOB) SRTSEQ(\*HEX)

| The optimizer can use index 1 (\*HEX) or 2 (\*LANGIDUNQ).

| **Example 12:** Grouping with unique-weight sequence table.

```
OPNQRYF FILE(STAFF) GRPFLD(JOB)
          SRTSEQ(*LANGIDUNQ) LANGID(ENU)
```

The optimizer can use index 1 (\*HEX) or 2 (\*LANGIDUNQ).

**Example 13:** Grouping with shared-weight sequence table.

```
OPNQRYF FILE(STAFF) GRPFLD(JOB)
          SRTSEQ(*LANGIDSHR) LANGID(ENU)
```

The optimizer can only use index 3 (\*LANGIDSHR).

### More Examples

In the following examples, the access paths (numbers 1, 2, and 3) from the examples 1 through 13 still exist over the JOB field.

In examples 14 through 20 there are access paths (numbers 4, 5, and 6) built over the JOB and SALARY fields. The access paths use the following sort sequence tables:

1. SRTSEQ(\*HEX)
2. SRTSEQ(\*LANGIDUNQ) LANGID(ENU)
3. SRTSEQ(\*LANGIDSHR) LANGID(ENU)
4. SRTSEQ(\*HEX)
5. SRTSEQ(\*LANGIDUNQ) LANGID(ENU)
6. SRTSEQ(\*LANGIDSHR) LANGID(ENU)

**Example 14:** Ordering and grouping on the same fields with a unique-weight sequence table.

```
OPNQRYF FILE(STAFF)
          SRTSEQ(*LANGIDUNQ) LANGID(ENU)
          GRPFLD(JOB SALARY) KEYFLD((JOB) (SALARY))
```

The optimizer can use index 5 (\*LANGIDUNQ) to satisfy both the grouping and ordering requirements. If index 5 did not exist, the optimizer would create an index using the \*LANGIDUNQ sequence table.

**Example 15:** Ordering and grouping on the same fields with ALWCPYDTA and a unique-weight sequence table.

```
OPNQRYF FILE(STAFF) ALWCPYDTA(*OPTIMIZE)
          SRTSEQ(*LANGIDUNQ) LANGID(ENU)
          GRPFLD(JOB SALARY) KEYFLD((JOB) (SALARY))
```

The optimizer can use index 5 (\*LANGIDUNQ) to satisfy both the grouping and ordering requirements. If index 5 does not exist, the optimizer would either:

- Create an index using a \*LANGIDUNQ sequence table.
- Use index 4 (\*HEX) to satisfy the grouping and perform a sort to satisfy the ordering.

**Example 16:** Ordering and grouping on the same fields with a shared-weight sequence table.

```
OPNQRYF FILE(STAFF)
          SRTSEQ(*LANGIDSHR) LANGID(ENU)
          GRPFLD(JOB SALARY) KEYFLD((JOB) (SALARY))
```

The optimizer can use index 6 (\*LANGIDSHR) to satisfy both the grouping and ordering requirements. If index 6 did not exist, the optimizer would create an index using a \*LANGIDSHR sequence table.

**Example 17:** Ordering and grouping on the same fields with ALWCPYDTA and a shared-weight sequence table.

```
OPNQRYF FILE(STAFF) ALWCPYDTA(*OPTIMIZE)
      SRTSEQ(*LANGIDSHR) LANGID(ENU)
      GRPFLD(JOB SALARY) KEYFLD((JOB)(SALARY))
```

The optimizer can use index 6 (\*LANGIDSHR) to satisfy both the grouping and ordering requirements. If index 6 did not exist, the optimizer would create an index using a \*LANGIDSHR sequence table.

**Example 18:** Ordering and grouping on different fields with a unique-weight sequence table.

```
OPNQRYF FILE(STAFF)
      SRTSEQ(*LANGIDUNQ) LANGID(ENU)
      GRPFLD(JOB SALARY) KEYFLD((SALARY)(JOB))
```

The optimizer can use index 4 (\*HEX) or 5 (\*LANGIDUNQ) to satisfy the grouping requirements. The grouping results are put into a temporary file. A temporary index using a \*LANGIDUNQ sequence table is built over the temporary result file to satisfy the ordering requirement.

**Example 19:** Ordering and grouping on different fields with ALWCPYDTA and a unique-weight sequence table.

```
OPNQRYF FILE(STAFF) ALWCPYDTA(*OPTIMIZE)
      SRTSEQ(*LANGIDUNQ) LANGID(ENU)
      GRPFLD(JOB SALARY) KEYFLD((SALARY)(JOB))
```

The optimizer can use index 4 (\*HEX) or 5 (\*LANGIDUNQ) to satisfy the grouping requirement. A sort then satisfies the ordering requirement.

**Example 20:** Ordering and grouping on different fields with ALWCPYDTA and a shared-weight sequence table.

```
OPNQRYF FILE(STAFF) ALWCPYDTA(*OPTIMIZE)
      SRTSEQ(*LANGIDSHR) LANGID(ENU)
      GRPFLD(JOB SALARY) KEYFLD((SALARY)(JOB))
```

The optimizer can use index 6 (\*LANGIDSHR) to satisfy the grouping requirement. A sort then satisfies the ordering requirement.

## Performance Comparisons with Other Database Functions

The Open Query File (OPNQRYF) command uses the same database support as logical files and join logical files. Therefore, the performance of functions like building a keyed access path or doing a join operation will be the same.

The selection functions done by the OPNQRYF command (for the QRYSLT and GRPSLT parameters) are similar to logical file select/omit. The main difference is that for the OPNQRYF command, the system decides whether to use access path selection or dynamic selection (similar to omitting or specifying the DYNLSLT

keyword in the DDS for a logical file), as a result of the access paths available on the system and what value was specified on the OPTIMIZE parameter.

## Considerations for Field Use

When the grouping function is used, all fields in the record format for the open query file (FORMAT parameter) and all key fields (KEYFLD parameter) must either be grouping fields (specified on the GRPFLD parameter) or mapped fields (specified on the MAPFLD parameter) that are defined using only grouping fields, constants, and aggregate functions. The aggregate functions are: %AVG, %COUNT, %MAX (using only one operand), %MIN (using only one operand), %STDDEV, %SUM, and %VAR. Group processing is required in the following cases:

- When you specify grouping field names on the GRPFLD parameter
- When you specify group selection values on the GRPSLT parameter
- When a mapped field that you specified on the MAPFLD parameter uses an aggregate function in its definition

Fields contained in a record format, identified on the FILE parameter, and defined (in the DDS used to create the file) with a usage value of N (neither input nor output) cannot be specified on any parameter of the OPNQRYF command. Only fields defined as either I (input-only) or B (both input and output) usage can be specified. Any fields with usage defined as N in the record format identified on the FORMAT parameter are ignored by the OPNQRYF command.

Fields in the open query file records normally have the same usage attribute (input-only or both input and output) as the fields in the record format identified on the FORMAT parameter, with the exceptions noted below. If the file is opened for any option (OPTION parameter) that includes output or update and any usage, and if any B (both input and output) field in the record format identified on the FORMAT parameter is changed to I (input only) in the open query file record format, then an information message is sent by the OPNQRYF command.

If you request join processing or group processing, or if you specify UNIQUEKEY processing, all fields in the query records are given input-only use. Any mapping from an input-only field from the file being processed (identified on the FILE parameter) is given input-only use in the open query file record format. Fields defined using the MAPFLD parameter are normally given input-only use in the open query file. A field defined on the MAPFLD parameter is given a value that matches the use of its constituent field if all of the following are true:

- Input-only is not required because of any of the conditions previously described in this section.
- The field-definition expression specified on the MAPFLD parameter is a field name (no operators or built-in functions).
- The field used in the field-definition expression exists in one of the file, member, or record formats specified on the FILE parameter (not in another field defined using the MAPFLD parameter).
- The base field and the mapped field are compatible field types (the mapping does not mix numeric and character field types, unless the mapping is between zoned and character fields of the same length).
- If the base field is binary with nonzero decimal precision, the mapped field must also be binary and have the same precision.

## Considerations for Files Shared in a Job

In order for your application program to use the open data path built by the Open Query File (OPNQRYF) command, your program must share the query file. If your program does not open the query file as shared, then it actually does a full open of the file it was originally compiled to use (not the query open data path built by the OPNQRYF command). Your program will share the query open data path, depending on the following conditions:

- Your application program must open the file as shared. Your program meets this condition when the first or only member queried (as specified on the FILE parameter) has an attribute of SHARE(\*YES). If the first or only member has an attribute of SHARE(\*NO), then you must specify SHARE(\*YES) in an Override with Database File (OVRDBF) command before calling your program.
- The file opened by your application program must have the same name as the file opened by the OPNQRYF command. Your program meets this condition when the file specified in your program has the same file and member name as the first or only member queried (as specified on the FILE parameter). If the first or only member has a different name, then you must specify an Override with Database File (OVRDBF) command of the name of the file your program was compiled against to the name of the first or only member queried.
- Your program must be running in the same activation group to which the query open data path (ODP) is scoped. If the query ODP is scoped to the job, your program may run in any activation group within the job.

The OPNQRYF command never shares an existing open data path in the job or activation group. A request to open a query file fails with an error message if the open data path has the same library, file, and member name that is in the open request, and if either of the following is true:

- OPNSCOPE(\*ACTGRPDFN) or OPNSCOPE(\*ACTGRP) is specified for the OPNQRYF command, and the open data path is scoped to the same activation group or job from which the OPNQRYF command is run.
- OPNSCOPE(\*JOB) is specified for the OPNQRYF command, and the open data path is scoped to the same job from which the OPNQRYF command is run.

Subsequent shared opens adhere to the same open options (such as SEQONLY) that were in effect when the OPNQRYF command was run.

See “Sharing Database Files in the Same Job or Activation Group” on page 5-8 for more information about sharing files in a job or activation group.

## Considerations for Checking If the Record Format Description Changed

If record format level checking is indicated, the format level number of the open query file record format (identified on the FORMAT parameter) is checked against the record format your program was compiled against. This occurs when your program shares the previously opened query file. Your program’s shared open is checked for record format level if the following conditions are met:

- The first or only file queried (as specified on the FILE parameter) must have the LVLCHK(\*YES) attribute.
- There must not be an override of the first or only file queried to LVLCHK(\*NO).

## Other Run Time Considerations

Overrides can change the name of the file, library, and member that should be processed by the open query file. (However, any parameter values other than TOFILE, MBR, LVLCHK, INHWRT, or SEQONLY specified on an Override with Database File (OVRDBF) command are ignored by the OPNQRYF command.) If a name change override applies to the first or only member queried, any additional overrides must be against the new name, not the name specified for the FILE parameter on the OPNQRYF command.

## Copying from an Open Query File

The Copy from Query File (CPYFRMQRYF) command can be used to copy from an open query file to another file or to print a formatted listing of the records. Any open query file, except those using distributed data management (DDM) files, specified with the input, update, or all operation value on the FILE parameter of the Open Query File (OPNQRYF) command can be copied using the CPYFRMQRYF command. The CPYFRMQRYF command cannot be used to copy to logical files. For more information, see the *Data Management Guide*.

Although the CPYFRMQRYF command uses the open data path of the open query file, it does not open the file. Consequently, you do not have to specify SHARE(\*YES) for the database file you are copying.

The following are examples of how the OPNQRYF and CPYFRMQRYF commands can be used.

### **Example 1:** Building a file with a subset of records

Assume you want to create a file from the CUSTOMER/ADDRESS file containing only records where the value of the STATE field is Texas. You can specify the following:

```
OPNQRYF FILE(CUSTOMER/ADDRESS) QRYSLT('STATE *EQ "TEXAS"')
CPYFRMQRYF FROMOPNID(ADDRESS) TOFILE(TEXAS/ADDRESS) CRTFILE(*YES)
```

### **Example 2:** Printing records based on selection

Assume you want to print all records from FILEA where the value of the CITY field is Chicago. You can specify the following:

```
OPNQRYF FILE(FILEA) QRYSLT('CITY *EQ "CHICAGO"')
CPYFRMQRYF FROMOPNID(FILEA) TOFILE(*PRINT)
```

### **Example 3:** Copying a subset of records to a diskette

Assume you want to copy all records from FILEB where the value of FIELDDB is 10 to a diskette. You can specify the following:

```
OPNQRYF FILE(FILEB) QRYSLT('FIELDDB *EQ "10"') OPNID(MYID)
CPYFRMQRYF FROMOPNID(MYID) TOFILE(DISK1)
```

### **Example 4:** Creating a copy of the output of a dynamic join

Assume you want to create a physical file that has the format and data of the join of FILEA and FILEB. Assume the files contain the following fields:

FILEA	FILEB	JOINAB
Cust	Cust	Cust
Name	Amt	Name
Addr		Amt

The join field is Cust, which exists in both files. To join the files and save a copy of the results in a new physical file MYLIB/FILEC, you can specify:

```
OPNQRYF FILE(FILEA FILEB) FORMAT(JOINAB) +
  JFLD((FILEA/CUST FILEB/CUST)) +
  MAPFLD((CUST 'FILEA/CUST')) OPNID(QRYFILE)
CPYFRMQRYF FROMOPNID(QRYFILE) TOFILE(MYLIB/FILEC) CRTFILE(*YES)
```

The file MYLIB/FILEC will be created by the CPYFRMQRYF command. The file will have file attributes like those of FILEA although some file attributes may be changed. The format of the file will be like JOINAB. The file will contain the data from the join of FILEA and FILEB using the Cust field. File FILEC in library MYLIB can be processed like any other physical file with CL commands, such as the Display Physical File Member (DSPPFM) command and utilities, such as Query. For more information about the CPYFRMQRYF command and other copy commands, see the *Data Management Guide*.

## Typical Errors When Using the Open Query File (OPNQRYF) Command

Several functions must be correctly specified for the OPNQRYF command and your program to get the correct results. The Display Job (DSPJOB) command is your most useful tool if problems occur. This command supports both the open files option and the file overrides option. You should look at both of these if you are having problems.

These are the most common problems and how to correct them:

- Shared open data path (ODP). The OPNQRYF command operates through a shared ODP. In order for the file to process correctly, the member must be opened for a shared ODP. If you are having problems, use the open files option on the DSPJOB command to determine if the member is opened and has a shared ODP.

There are normally two reasons that the file is not open:

- The member to be processed must be SHARE(\*YES). Either use an Override with Database File (OVRDBF) command or permanently change the member.
- The file is closed. You have run the OPNQRYF command with the OPNSCOPE(\*ACTGRPDFN) or TYPE(\*NORMAL) parameter option from a program that was running in the default activation group at a higher level in the call stack than the program that is getting an error message or that is simply running the Reclaim Resources (RCLRSC) command. This closes the open query file because it was opened from a program at a higher level in the call stack than the program that ran the RCLRSC command. If the open query file was closed, you must run the OPNQRYF command again. Note that when using the OPNQRYF command with the TYPE(\*NORMAL) parameter option on releases prior to Version 2 Release 3, the open query file is closed even if it was opened from the same program that reclaims the resources.

- Level check. Level checking is normally used because it ensures that your program is running against the same record format that the program was com-

piled with. If you are experiencing level check problems, it is normally because of one of the following:

- The record format was changed since the program was created. Creating the program again should correct the problem.
- An override is directing the program to an incorrect file. Use the file overrides option on the DSPJOB command to ensure that the overrides are correctly specified.
- The FORMAT parameter is needed but is either not specified or incorrectly specified. When a file is processed with the FORMAT parameter, you must ensure:
  - The Override with Database File (OVRDBF) command, used with the TOFILE parameter, describes the first file on the FILE parameter of the Open Query File (OPNQRYF) command.
  - The FORMAT parameter identifies the file that contains the format used to create the program.
- The FORMAT parameter is used to process a format from a different file (for example, for group processing), but SHARE(\*YES) was not requested on the OVRDBF command.
- The file to be processed is at end of file. The normal use of the OPNQRYF command is to process a file sequentially where you can only process the file once. At that point, the position of the file is at the end of the file and you will not receive any records if you attempt to process it again. To process the file again from the start, you must either run the OPNQRYF command again or reposition the file before processing. You can reposition the file by using the Position Database File (POSDBF) command, or through a high-level language program statement.
- No records exist. This can be caused when you use the FORMAT keyword, but do not specify the OVRDBF command.
- Syntax errors. The system found an error in the specification of the OPNQRYF command.
- Operation not valid. The definition of the query does not include the KEYFLD parameter, but the high-level language program attempts to read the query file using a key field.
- Get option not valid. The high-level language program attempted to read a record or set a record position before the current record position, and the query file used either the group by option, the unique key option, or the distinct option on the SQL/400 statement.



## Chapter 7. Basic Database File Operations

The basic database file operations that can be performed in a program are discussed in this chapter. The operations include: setting a position in the database file, reading records from the file, updating records in the file, adding records to the file, and deleting records from the file.

### Setting a Position in the File


After a file is opened by a job, the system maintains a position in the file for that job. The file position is used in processing the file. For example, if a program does a read operation requesting the next sequential record, the system uses the file position to determine which record to return to the program. The system will then set the file position to the record just read, so that another read operation requesting the next sequential record will return the correct record. The system keeps track of all file positions for each job. In addition, each job can have multiple positions in the same file.

The file position is first set to the position specified in the POSITION parameter on the Override with Database File (OVRDBF) command. If you do not use an OVRDBF command, or if you take the default for the POSITION parameter, the file position is set just before the first record in the member's access path.

A program can change the current file position by using the appropriate high-level language program file positioning operation (for example, SETLL in the RPG/400 language or START in the COBOL/400 language). A program can also change the file position by using the CL Position Database File (POSDBF) command.

**Note:** File positioning by means of the Override with Database File (OVRDBF) command does not occur until the next time the file is opened. Because a file can be opened only once within a CL program, this command cannot be used within a single CL program to affect what will be read through the RCVF command.

At end of file, after the last read, the file member is positioned to \*START or \*END file position, depending on whether the program was reading forward or backward through the file. The following diagram shows \*START and \*END file positions.



RV2F649-0

Only a read operation, force-end-of-data operation, high-level language positioning operation, or specific CL command to change the file position can change the file position. Add, update, and delete operations do not change the file position. After

a read operation, the file is positioned to the new record. This record is then returned to your program. After the read operation is completed, the file is positioned at the record just returned to your program. If the member is open for input, a force-end-of-data operation positions the file after the last record in the file (\*END) and sends the end-of-file message to your program.

For sequential read operations, the current file position is used to locate the next or previous record on the access path. For read-by-key or read-by-relative-record-number operations, the file position is not used. If POSITION(\*NONE) is specified at open time, no starting file position is set. In this case, you must establish a file position in your program, if you are going to read sequentially.

If end-of-file delay was specified for the file on an Override With Database File (OVRDBF) command, the file is not positioned to \*START or \*END when the program reads the last record. The file remains positioned at the last record read. A file with end-of-file delay processing specified is positioned to \*START or \*END only when a force-end-of-data (FEOD) occurs or a controlled job end occurs. For more information about end-of-file delay, see "Waiting for More Records When End of File Is Reached" on page 7-5.

You can also use the Position Database File (POSDBF) command to set or change the current position in your file for files opened using either the Open Database File (OPNDBF) command or the Open Query File (OPNQRYF) command.

---

## Reading Database Records

The AS/400 system provides a number of ways to read database records. The next sections describe those ways in detail. (Some high-level languages do not support all of the read operations available on the system. See your high-level language guide for more information about reading database records.)

### Reading Database Records Using an Arrival Sequence Access Path

The system performs the following read operations based on the operations you specify using your high-level language. These operations are allowed if the file was defined with an arrival sequence access path; or if the file was defined with a keyed sequence access path with the ignore-keyed-sequence-access-path option specified in the program, on the Open Database File (OPNDBF) command, or on the Open Query File (OPNQRYF) command. See "Ignoring the Keyed Sequence Access Path" on page 5-4 for more details about the option to ignore a keyed sequence access path.

**Note:** Your high-level language may not allow all of the following read operations. Refer to your high-level language guide to determine which operations are allowed by the language.

**Read Next:** Positions the file to and gets the next record that is not deleted in the arrival sequence access path. Deleted records between the current position in the file and the next active record are skipped. (The READ statement in the RPG/400 language and the READ NEXT statement in the COBOL/400 language are examples of this operation.)

**Read Previous:** Positions the file to and gets the previous active record in the arrival sequence access path. Deleted records between the current file position and the previous active record are skipped. (The READP statement in the

RPG/400 language and the READ PRIOR statement in the COBOL/400 language are examples of this operation.)

**Read First:** Positions the file to and gets the first active record in the arrival sequence access path.

**Read Last:** Positions the file to and gets the last active record in the arrival sequence access path.

**Read Same:** Gets the record that is identified by the current position in the file. The file position is not changed.

**Read by Relative Record Number:** Positions the file to and gets the record in the arrival sequence access path that is identified by the relative record number. The relative record number must identify an active record and must be less than or equal to the largest active relative record number in the member. This operation also reads the record in the arrival sequence access path identified by the current file position plus or minus a specified number of records. (The CHAIN statement in the RPG/400 language and the READ statement in the COBOL/400 language are examples of this operation.) Special consideration should be given to creating or changing a file to reuse deleted records if the file is processed by relative record processing. For more information, see “Reusing Deleted Records” on page 5-3.

## Reading Database Records Using a Keyed Sequence Access Path

The system performs the following read operations based on the statements you specify using your high-level language. These operations can be used with a keyed sequence access path to get database records.

When a keyed sequence access path is used, a read operation cannot position to the storage occupied by a deleted record.

**Note:** Your high-level language may not allow all of the following operations. Refer to your high-level language guide to determine which operations are allowed by the language.

**Read Next:** Gets the next record on the keyed sequence access path. If a record format name is specified, this operation gets the next record in the keyed sequence access path that matches the record format. The current position in the file is used to locate the next record. (The READ statement in the RPG/400 language and the READ NEXT statement in the COBOL/400 language are examples of this operation.)

**Read Previous:** Gets the previous record on the keyed sequence access path. If a record format name is specified, this operation gets the previous record on the keyed sequence access path that matches the record format. The current position in the file is used to locate the previous record. (The READP statement in the RPG/400 language and the READ PRIOR statement in the COBOL/400 language are examples of this operation.)

**Read First:** Gets the first record on the keyed sequence access path. If a record format name is specified, this operation gets the first record on the access path with the specified format name.

**Read Last:** Gets the last record on the keyed sequence access path. If a record format name is specified, this operation gets the last record on the access path with the specified format name.

**Read Same:** Gets the record that is identified by the current file position. The position in the file is not changed.

**Read by Key:** Gets the record identified by the key value. Key operations of equal, equal or after, equal or before, read previous key equal, read next key equal, after, or before can be specified. If a format name is specified, the system searches for a record of the specified key value and record format name. If a format name is not specified, the entire keyed sequence access path is searched for the specified key value. If the key definition for the file includes multiple key fields, a partial key can be specified (you can specify either the number of key fields or the key length to be used). This allows you to do generic key searches. If the program does not specify a number of key fields, the system assumes a default number of key fields. This default varies depending on if a record format name is passed by the program. If a record format name is passed, the default number of key fields is the total number of key fields defined for that format. If a record format name is not passed, the default number of key fields is the maximum number of key fields that are common across all record formats in the access path. The program must supply enough key data to match the number of key fields assumed by the system. (The CHAIN statement in the RPG/400 language and the READ statement in the COBOL/400 language are examples of this operation.)

**Read by Relative Record Number:** For a keyed sequence access path, the relative record number can be used. This is the relative record number in the arrival sequence, even though the member opened has a keyed sequence access path. If the member contains multiple record formats, a record format name must be specified. In this case, you are requesting a record in the associated physical file member that matches the record format specified. If the member opened contains select/omit statements and the record identified by the relative record number is omitted from the keyed sequence access path, an error message is sent to your program and the operation is not allowed. After the operation is completed, the file is positioned to the key value in the keyed sequence access path that is contained in the physical record, which was identified by the relative record number. This operation also gets the record in the keyed sequence access path identified by the current file position plus or minus some number of records. (The CHAIN statement in the RPG/400 language and the READ statement in the COBOL/400 language are examples of this operation.)

**Read when Logical File Shares an Access Path with More Keys:** When the FIFO, LIFO, or FCFO keyword is not specified in the data description specifications (DDS) for a logical file, the logical file can implicitly share an access path that has more keys than the logical file being created. This sharing of a partial set of keys from an existing access path can lead to perceived problems for database read operations that use these partially shared keyed sequence access paths. The problems will appear to be:

- Records that should be read, are never returned to your program
- Records are returned to your program multiple times

What is actually happening is that your program or another currently active program is updating the physical file fields that are keys within the partially shared keyed sequence access path, but that are not actual keys for the logical file that is being

used by your program (the fields being updated are beyond the number of keys known to the logical file being used by your program). The updating of the actual key fields for a logical file by your program or another program has always yielded the above results. The difference with partially shared keyed sequence access paths is that the updating of the physical file fields that are keys beyond the number of keys known to the logical file can cause the same consequences.

If these consequences caused by partially shared keyed sequence access paths are not acceptable, the FIFO, LIFO, or FCFO keyword can be added to the DDS for the logical file, and the logical file created again.

## Waiting for More Records When End of File Is Reached

End-of-file delay is a method of continuing to read sequentially from a database file (logical or physical) after an end-of-file condition occurs. When an end-of-file condition occurs on a file being read sequentially (for example, next/previous record) and you have specified an end-of-file delay time (EOFDLY parameter on the Override with Database File [OVRDBF] command), the system waits for the time you specified. At the end of the delay time, another read is done to determine if any new records were added to the file. If records were added, normal record processing is done until an end-of-file condition occurs again. If records were not added to the file, the system waits again for the time specified. Special consideration should be taken when using end-of-file delay on a logical file with select/omit specifications, opened so that the keyed sequence access path is not used. In this case, once end-of-file is reached, the system retrieves only those records added to a based-on physical file that meet the select/omit specifications of the logical file.

Also, special consideration should be taken when using end-of-file delay on a file with a keyed sequence access path, opened so that the keyed sequence access path is used. In this case, once end-of-file is reached, the system retrieves only those records added to the file or those records updated in the file that meet the specification of the read operation using the keyed sequence access path.

For example, end-of-file delay is used on a keyed file that has a numeric key field in ascending order. An application program reads the records in the file using the keyed sequence access path. The application program performs a read next operation and gets a record that has a key value of 99. The application program performs another read next and no more records are found in the file, so the system attempts to read the file again after the specified end-of-file delay time. If a record is added to the file or a record is updated, and the record has a key value less than 99, the system does not retrieve the record. If a record is added to the file or a record is updated and the record has a key value greater than or equal to 99, the system retrieves the record.

For end-of-file delay times equal to or greater than 10 seconds, the job is eligible to be removed from main storage during the wait time. If you do not want the job eligible to be moved from main storage, specify PURGE(\*NO) on the Create Class (CRTCLS) command for the CLASS the job is using.

To indicate which jobs have an end-of-file delay in progress, the status field of the Work with Active Jobs (WRKACTJOB) display shows an end-of-file wait or end-of-file activity level for jobs that are waiting for a record.

If a job uses end-of-file-delay and commitment control, it can hold its record locks for a longer period of time. This increases the chances that some other job can try

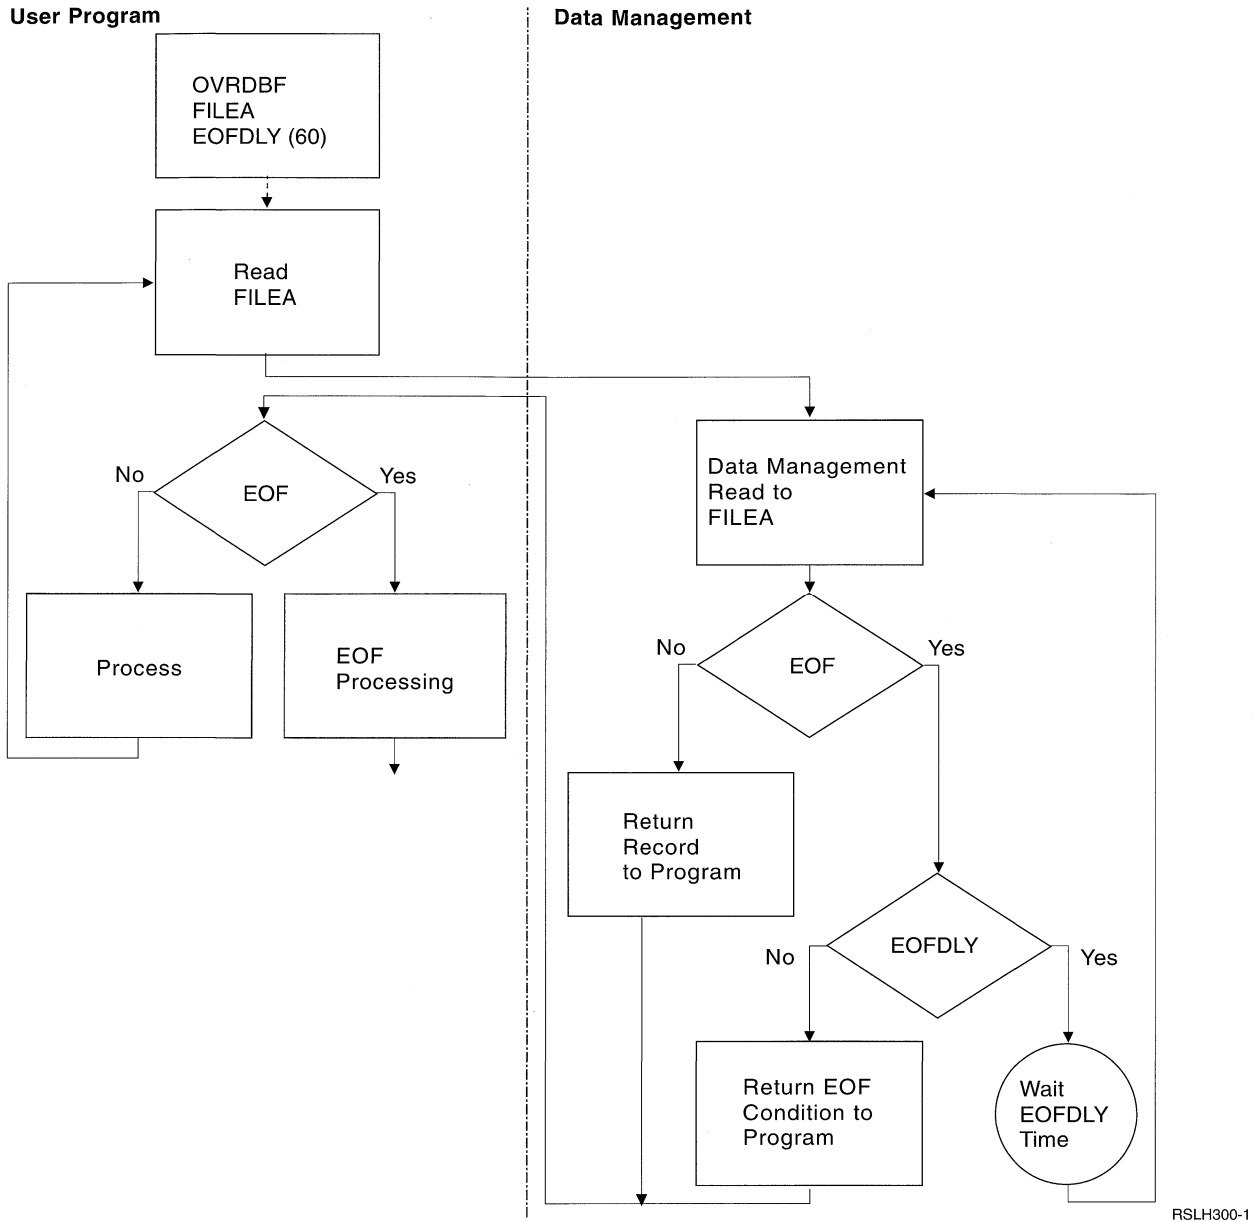
to access those same records and be locked out. For that reason, be careful when using end-of-file-delay and commitment control in the same job.

If a file is shared, the Override with Database File (OVRDBF) command specifying an end-of-file delay must be requested before the *first* open of the file because overrides are ignored that are specified after the shared file is opened.

There are several ways to end a job that is waiting for more records because of an end-of-file-delay specified on the Override with Database File (OVRDBF) command:

- Write a record to the file with the end-of-file-delay that will be recognized by the application program as a last record. The application program can then specify a force-end-of-data (FEOD) operation. An FEOD operation allows the program to complete normal end-of-file processing.
- Do a controlled end of a job by specifying OPTION(\*CNTRLD) on the End Job (ENDJOB) command, with a DELAY parameter value time greater than the EOFDLY time. The DELAY parameter time specified must allow time for the EOFDLY time to run out, time to process any new records that have been put in the file, and any end-of-file processing required in your application. After new records are processed, the system signals end of file, and a normal end-of-file condition occurs.
- Specify OPTION(\*IMMED) on the End Job (ENDJOB) command. No end-of-file processing is done.
- If the job is interactive, press the System Request key to end the last request.

The following is an example of end-of-file delay operation:



The actual processing of the EOFDLY parameter is more complex than shown because it is possible to force a true end-of-file if `OPTION(*CNTRLD)` on the End Job (`ENDJOB`) command is used with a long delay time.

The job does not become active whenever a new record is added to the file. The job becomes active after the specified end-of-file delay time ends. When the job becomes active, the system checks for any new records. If new records were added, the application program gets control and processes all new records, then waits again. Because of this, the job takes on the characteristic of a batch job when it is processing. For example, it normally processes a batch of requests. When the batch is completed, the job becomes inactive. If the delay is small, you can cause excessive system overhead because of the internal processing required to start the job and check for new records. Normally, only a small amount of overhead is used for a job waiting during end-of-file delay.

**Note:** When the job is inactive (waiting) it is in a long-wait status, which means it was released from an activity level. After the long-wait status is satisfied, the system reschedules the job in an activity level. (See the *Work Management Guide* for more information about activity levels.)

## Releasing Locked Records

The system automatically releases a locked record when the record is updated, deleted, or when you read another record in the file. However, you may want to release a locked record without performing these operations. Some high-level languages support an operation to release a locked record. See your high-level language guide for more information about releasing record locks.

**Note:** The rules for locking are different if your job is running under commitment control. See the *Advanced Backup and Recovery Guide* for more details.

---

## Updating Database Records

The update operation allows you to change an existing database record in a logical or physical file. (The UPDAT statement in the RPG/400 language and the REWRITE statement in the COBOL/400 language are examples of this type operation.) Before you update a database record, the record must first be read and locked. The lock is obtained by specifying the update option on any of the read operations listed under the “Reading Database Records Using an Arrival Sequence Access Path” on page 7-2 or “Reading Database Records Using a Keyed Sequence Access Path” on page 7-3.

If you issue several read operations with the update option specified, each read operation releases the lock on the previous record before attempting to locate and lock the new record. When you do the update operation, the system assumes that you are updating the currently locked record. Therefore, you do not have to identify the record to be updated on the update operation. After the update operation is done, the system releases the lock.

**Note:** The rules for locking are different if your job is running under commitment control. See the *Advanced Backup and Recovery Guide* for more details.

If the update operation changes a key field in an access path for which immediate maintenance is specified, the access path is updated if the high-level language allows it. (Some high-level languages do not allow changes to the key field in an update operation.)

If you request a read operation on a record that is already locked for update and if your job is running under a commitment control level of \*ALL or \*CS (cursor stability), then you must wait until the record is released or the time specified by the WAITRCD parameter on the create file or override commands has been exceeded. If the WAITRCD time is exceeded without the lock being released, an exception is returned to your program and a message is sent to your job stating the file, member, relative record number, and the job which has the lock. If the job that is reading records is not running under a commitment control level of \*ALL or \*CS, the job is able to read a record that is locked for update.



---

## Adding Database Records

The write operation is used to add a new record to a physical database file member. (The WRITE statement in the RPG/400 language and the WRITE statement in the COBOL/400 language are examples of this operation.) New records can be added to a physical file member or to a logical file member that is based on the physical file member. When using a multiple format logical file, a record format name must be supplied to tell the system which physical file member to add the record to.

The new record is normally added at the end of the physical file member. The next available relative record number (including deleted records) is assigned to the new record. Some high-level languages allow you to write a new record over a deleted record position (for example, the WRITE statement in COBOL/400 when the file organization is defined as RELATIVE). For more information about writing records over deleted record positions, see your high-level language guide.

If the physical file to which records are added reuses deleted records, the system tries to insert the records into slots that held deleted records. Before you create or change a file to reuse deleted records, you should review the restrictions and tips for use to determine whether the file is a candidate for reuse of deleted record space. For more information on reusing deleted record space, see “Reusing Deleted Records” on page 5-3.

If you are adding new records to a file member that has a keyed access path, the new record appears in the keyed sequence access path immediately at the location defined by the record key. If you are adding records to a logical member that contains select/omit values, the omit values can prevent the new record from appearing in the member’s access path.

The SIZE parameter on the Create Physical File (CRTPF) and Create Source Physical File (CRTSRCPF) commands determines how many records can be added to a physical file member.

## Identifying Which Record Format to Add in a File with Multiple Formats

If your application uses a file name instead of a record format name for records to be added to the database, and if the file used is a logical file with more than one record format, you need to write a format selector program to determine where a record should be placed in the database. A format selector can be a CL program or a high-level language program.

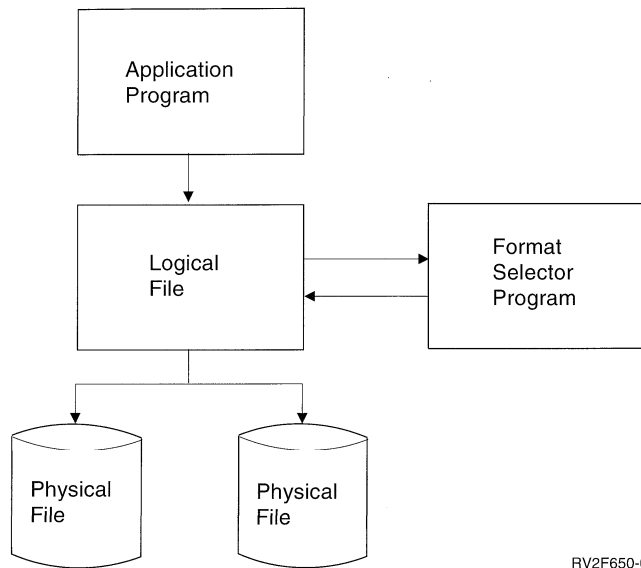
A format selector program must be used if all of the following are true:

- The logical file is not a join and not a view logical file.
- The logical file is based on multiple physical files.
- The program uses a file name instead of a record format name on the add operation.

If you do not write a format selector program for this situation, your program ends with an error when it tries to add a record to the database.

**Note:** A format selector program cannot be used to select a member if a file has multiple members; it can only select a record format.

When an application program wants to add a record to the database file, the system calls the format selector program. The format selector program examines the record and specifies the record format to be used. The system then adds the record to the database file using the specified record format name.



The following example shows the programming statements for a format selector program written in the RPG/400 language:

```

CL0N01N02N03Factor1+++0pcdeFactor2+++ResultLenDHHiLoEqComments+++...
++++
C          *ENTRY    PLIST
C          PARM      RECORD 80
C* The length of field RECORD must equal the length of
C* the longest record expected.
C          PARM      FORMAT 10
C          MOVELRECORD  BYTE  1
C          BYTE     IFEQ 'A'
C          MOVEL 'HDR'  FORMAT
C          ELSE
C          MOVEL 'DTL'  FORMAT
C          END
  
```

The format selector receives the record in the first parameter; therefore, this field must be declared to be the length of the longest record expected by the format selector. The format selector can access any portion of the record to determine the record format name. In this example, the format selector checks the first character in the record for the character A. If the first character is A, the format selector moves the record format name HDR into the second parameter (FORMAT). If the character is not A, the format selector moves the record format name DTL into the second parameter.

The format selector uses the second parameter, which is a 10-character field, to pass the record format name to the system. When the system knows the name of the record format, it adds the record to the database.

You do not need a format selector if:

- You are doing updates only. For updates, your program already retrieved the record, and the system knows which physical file the record came from.
- Your application program specifies the record format name instead of a file name for an add or delete operation.
- All the records used by your application program are contained in one physical file.

To create the format selector, you use the create program command for the language in which you wrote the program. You cannot specify `USRPRF(*OWNER)` on the create command. The format selector must run under the user's user profile not the owner's user profile.

In addition, for security and integrity and because performance would be severely affected, you must not have any calls or input/output operations within the format selector.

The name of the format selector is specified on the `FMTSLR` parameter of the Create Logical File (`CRTL`), Change Logical File (`CHGLF`), or Override with Database File (`OVRDBF`) command. The format selector program does not have to exist when the file is created, but it must exist when the application program is run.

## Using the Force-End-Of-Data Operation

The force-end-of-data (FEOD) operation allows you to force all changes to a file made by your program to auxiliary storage. Normally, the system determines when to force changes to auxiliary storage. However, you can use the FEOD operation to ensure that all changes are forced to auxiliary storage.

The force-end-of-data (FEOD) operation also allows you to position to either the beginning or the end of a file if the file is open for input operations. `*START` sets the beginning or starting position in the database file member currently open to just before the first record in the member (the first sequential read operation reads the first record in the current member). If `MBR(*ALL)` processing is in effect for the override with Database File (`OVRDBF`) command, a read previous operation gets the last record in the previous member. If a read previous operation is done and the previous member does not exist, the end of file message (CPF5001) is sent. `*END` sets the position in the database file member currently open to just after the last record in the member (a read previous operation reads the last record in the current member). If `MBR(*ALL)` processing is in effect for the Override with Database File (`OVRDBF`) command, a read next operation gets the first record in the next member. If a read next operation is done and the next member does not exist, the end of file message (CPF5001) occurs.

See your high-level language guide for more information about the FEOD operation (some high-level languages do not support the FEOD operation).

---

## Deleting Database Records

The delete operation allows you to delete an existing database record. (The `DELET` statement in the RPG/400 language and the `DELETE` statement in the COBOL/400 language are examples of this operation.) To delete a database record, the record must first be read and locked. The record is locked by specifying the update option on any of the read operations listed under "Reading Database Records Using an Arrival Sequence Access Path" on page 7-2 or "Reading

Database Records Using a Keyed Sequence Access Path” on page 7-3. The rules for locking records for deletion and identifying which record to delete are the same as for update operations.

**Note:** Some high-level languages do *not* require that you read the record first. These languages allow you to simply specify which record you want deleted on the delete statement. For example, the RPG/400 language allows you to delete a record without first reading it.

When a database record is deleted, the physical record is marked as deleted. This is true even if the delete operation is done through a logical file. A deleted record *cannot* be read. The record is removed from all keyed sequence access paths that contain the record. The relative record number of the deleted record remains the same. All other relative record numbers within the physical file member do not change.

The space used by the deleted record remains in the file, but it is not reused until:

- The Reorganize Physical File Member (RGZPFM) command is run to compress and free these spaces in the file member. See “Reorganizing Data in Physical File Members” on page 10-3 for more information about this command.
- Your program writes a record to the file by relative record number and the relative record number used is the same as that of the deleted record.

**Note:** The system tries to reuse deleted record space automatically if the file has the reuse deleted record space attribute specified. For more information, see “Reusing Deleted Records” on page 5-3.

The system does not allow you to retrieve the data for a deleted record. You can, however, write a new record to the position (relative record number) associated with a deleted record. The write operation replaces the deleted record with a new record. See your high-level language guide for more details about how to write a record to a specific position (relative record number) in the file.

To write a record to the relative record number of a deleted record, that relative record number must exist in the physical file member. You can delete a record in the file using the delete operation in your high-level language. You can also delete records in your file using the Initialize Physical File Member (INZPFM) command. The INZPFM command can initialize the entire physical file member to deleted records. For more information about the INZPFM command, see “Initializing Data in a Physical File Member” on page 10-2.

---

## Chapter 8. Closing a Database File

When your program completes processing a database file member, it should close the file. Closing a database file disconnects your program from the file. The close operation releases all record locks and releases all file member locks, forces all changes made through the open data path (ODP) to auxiliary storage, then destroys the ODP. (When a shared file is closed but the ODP remains open, the functions differ. For more information about shared files, see “Sharing Database Files in the Same Job or Activation Group” on page 5-8.)

The ways you can close a database file in a program include:

- High-level language close statements
- Close File (CLOF) command
- Reclaim Resources (RCLRSC) command

Most high-level languages allow you to specify that you want to close your database files. For more information about how to close a database file in a high-level language program, see your high-level language guide.

You can use the Close File (CLOF) command to close database files that were opened using either the Open Database File (OPNDBF) or Open Query File (OPNQRYF) commands.

You can also close database files by running the Reclaim Resources (RCLRSC) command. The RCLRSC command releases all locks (except, under commitment control, locks on records that were changed but not yet committed), forces all changes to auxiliary storage, then destroys the open data path for that file. You can use the RCLRSC command to allow a calling program to close a called program's files. (For example, if the called program returns to the calling program without closing its files, the calling program can then close the called program's files.) However, the normal way of closing files in a program is with the high-level language close operation or through the Close File (CLOF) command. For more information on resource reclamation in the integrated language environment see Appendix A in *Integrated Language Environment Concepts*, SC09-1524.

If a job ends normally (for example, a user signs off) and all the files associated with that job were not closed, the system automatically closes all the remaining open files associated with that job, forces all changes to auxiliary storage, and releases all record locks for those files. If a job ends abnormally, the system also closes all files associated with that job, releases all record locks for those files, and forces all changes to auxiliary storage.



---

## Chapter 9. Handling Database File Errors in a Program

Error conditions detected during processing of a database file cause messages to be sent to the program message queue for the program processing the file or cause an inquiry message to be sent to the system operator message queue. In addition, file errors and diagnostic information generally appear to your program as return codes and status information in the file feedback area. (For example, the COBOL/400 language sets a return code in the file status field, if it is defined in the program.) For more information about handling file errors in your program, see your high-level language guide.

If your programming language allows you to monitor for error messages, you can choose which ones you wish to monitor for. The following messages are a small sample of the error messages you can monitor (see your high-level language guide and the *CL Reference* manual for a complete list of errors and messages you can monitor):

---

Message Identifier	Description
CPF5001	End of file reached
CPF5006	Record not found
CPF5007	Record deleted
CPF5018	Maximum file size reached
CPF5025	Read attempted past *START or *END
CPF5026	Duplicate key
CPF5027	Record in use by another job
CPF5028	Record key changed
CPF5029	Data mapping error
CPF5030	Partial damage on member
CPF5031	Maximum number of record locks exceeded
CPF5032	Record already allocated to job
CPF5033	Select/omit error
CPF5034	Duplicate key in another member's access path
CPF5040	Omitted record not retrieved
CPF5072	Join value in member changed
CPF5079	Commitment control resource limit exceeded
CPF5084	Duplicate key for uncommitted key
CPF5085	Duplicate key for uncommitted key in another access path
CPF5090	Unique access path problem prevents access to member
CPF5097	Key mapping error

---

**Note:** To display the full description of these messages, use the Display Message Description (DSPMSGD) command.

If you do not monitor for any of these messages, the system handles the error. The system also sets the appropriate error return code in the program. Depending on the error, the system can end the job or send a message to the operator requesting further action.

If a message is sent to your program while processing a database file member, the position in the file is not lost. It remains at the record it was positioned to before the message was sent, except:

- After an end-of-file condition is reached and a message is sent to the program, the file is positioned at \*START or \*END.
- After a conversion mapping message on a read operation, the file is positioned to the record containing the data that caused the message.



---

## Part 3. Managing Database Files

The chapters in this part include information on managing database files. This includes: adding new members, changing attributes of existing members, renaming members, or removing them from a database file. Information on member operations unique to physical files including initializing data, clearing data, reorganizing data, and displaying data in a physical file member is also included.

This section includes information on changing database file descriptions and attributes (including the effects of changing fields in file descriptions), changing physical file descriptions and attributes, changing logical file descriptions and attributes, and using database attributes and cross reference information is included.

Another topic covered is displaying information about the database files such as attributes, descriptions of fields, relationships between the fields, files used by programs, system cross reference files, and information on how to write a command output directly to a database file.

Also included is information to help you plan for recovery of your database files in the event of a system failure. This includes saving and restoring, journaling, using auxiliary storage, and commitment control. This section also includes information on access path recovery which includes rebuilding and journaling access paths.

A section on source files discusses source file concepts and reasons you would use a source file. Information on how to set up a source file, how to enter data into a source file, and ways to use a source file to create another object on the system is included.



---

## Chapter 10. Managing Database Members

Before you perform any input or output operations on a file, the file must have at least one member. As a general rule, database files have only one member, the one created when the file is created. The name of this member is the same as the file name, unless you give it a different name. Because most operations on database files assume that the member being used is the first member in the file, and because most files only have one member, you do not normally have to be concerned with, or specify, member names.

If a file contains more than one member, each member serves as a subset of the data in the file. This allows you to classify data easier. For example, you define an accounts receivable file. You decide that you want to keep data for a year in that file, but you frequently want to process data just one month at a time. For example, you create a physical file with 12 members, one named for each month. Then, you process each month's data separately (by individual member). You can also process several or all members together.

---

### Member Operations Common to All Database Files

The system supplies a way for you to:

- Add new members to an existing file.
- Change some attributes for an existing member (for example, the text describing the member) without having to re-create the member.
- Rename a member.
- Remove the member from the file.

The following section discusses these operations.

### Adding Members to Files

You can add members to files in any of these ways:

- Automatically. When a file is created using the Create Physical File (CRTPF) or Create Logical File (CRTLF) commands, the default is to automatically add a member (with the same name as the file) to the newly created file. (The default for the Create Source Physical File [CRTSRCPF] command is *not* to add a member to the newly created file.) You can specify a different member name using the MBR parameter on the create database file commands. If you do not want a member added when the file is created, specify \*NONE on the MBR parameter.
- Specifically. After the file is created, you can add a member using the Add Physical File Member (ADDPFM) or Add Logical File Member (ADDLFM) commands.
- Copy File (CPYF) command. If the member you are copying does not exist in the file being copied to, the member is added to the file by the CPYF command.

## Changing Member Attributes

You can use the Change Physical File Member (CHGPFM) or Change Logical File Member (CHGLFM) command to change certain attributes of a physical or a logical file member. For a physical file member, you can change the following parameters: SRCTYPE (the member's source type), EXPDATE (the member's expiration date), SHARE (whether the member can be shared within a job), and TEXT (the text description of the member). For a logical file member you can change the SHARE and TEXT parameters.

**Note:** You can use the Change Physical File (CHGPF) and Change Logical File (CHGLF) commands to change many other file attributes. For example, to change the maximum size allowed for each member in the file, you would use the SIZE parameter on the CHGPF command.

## Renaming Members

The Rename Member (RNMM) command changes the name of an existing member in a physical or logical file. The file name is not changed.

## Removing Members from Files

The Remove Member (RMVM) command is used to remove the member and its contents. Both the member data and the member itself are removed. After the member is removed, it can no longer be used by the system. This is different from just clearing or deleting the data from the member. If the member still exists, programs can continue to use (for example, add data to) the member.

---

## Physical File Member Operations

The following section describes member operations that are unique to physical file members. Those operations include initializing data, clearing data, reorganizing data, and displaying data in a physical file member.

### Initializing Data in a Physical File Member

To use relative record processing in a program, the database file must contain a number of record positions equal to the highest relative record number used in the program. Programs using relative-record-number processing sometimes require that these records be initialized.

You can use the Initialize Physical File Member (INZPFM) command to initialize members with one of two types of records:

- Default records
- Deleted records

You specify which type of record you want using the RECORDS parameter on the Initialize Physical File Member (INZPFM) command.

If you initialize records using default records, the fields in each new record are initialized to the default field values defined when the file was created. If no default field value was defined, then numeric fields are filled with zeros and character fields are filled with blanks.

Variable-length character fields have a zero-length default value. The default value for null-capable fields is the null value. The default value for dates, times, and

timestamps is the current date, time, or timestamp if no default value is defined. Program-described files have a default value of all blanks.

**Note:** You cannot initialize a default record if the UNIQUE keyword is specified in DDS for the physical file member or any associated logical file members. Otherwise, you would create a series of duplicate key records.

If the records are initialized to the default records, you can read a record by relative record number and change the data.

If the records were initialized to deleted records, you can change the data by adding a record using a relative record number of one of the deleted records. (You cannot add a record using a relative record number that was not deleted.)

Deleted records cannot be read; they only hold a place in the member. A deleted record can be changed by writing a new record over the deleted record. Refer to “Deleting Database Records” on page 7-11 for more information about processing deleted records.

## Clearing Data from Physical File Members

The Clear Physical File Member (CLRPFM) command is used to remove the data from a physical file member. After the clear operation is complete, the member description remains, but the data is gone.

## Reorganizing Data in Physical File Members

You can use the Reorganize Physical File Member (RGZPFM) command to:

- Remove deleted records to make the space occupied by them available for more records.
- Reorganize the records of a file in the order in which you normally access them sequentially, thereby minimizing the time required to retrieve records. This is done using the KEYFILE parameter. This may be advantageous for files that are primarily accessed in an order other than arrival sequence. A member can be reorganized using either of the following:
  - Key fields of the physical file
  - Key fields of a logical file based on the physical file
- Reorganize a source file member, insert new source sequence numbers, and reset the source date fields (using the SRCOPT and SRCSEQ parameters on the Reorganize Physical File Member command).
- Reclaim space in the variable portion of the file that was previously used by variable-length fields in the physical file format and that has now become fragmented.

For example, the following Reorganize Physical File Member (RGZPFM) command reorganizes the first member of a physical file using an access path from a logical file:

```
RGZPFM FILE(DSTPRODLB/ORDHDRP)
        KEYFILE(DSTPRODLB/ORDFILL ORDFILL)
```

The physical file ORDHDRP has an arrival sequence access path. It was reorganized using the access path in the logical file ORDFILL. Assume the key field is the *Order* field. The following illustrates how the records were arranged.

The following is an example of the original ORDHDRP file. Note that record 3 was deleted before the RGZPFM command was run:

Relative Record Number	Cust	Order	Ordate. . .
1	41394	41882	072480. . .
2	28674	32133	060280. . .
3	deleted	record	
4	56325	38694	062780. . .

The following example shows the ORDHDRP file reorganized using the *Order* field as the key field in ascending sequence:

Relative Record Number	Cust	Order	Ordate. . .
1	28674	32133	060280. . .
2	56325	38694	062780. . .
3	41394	41882	072480. . .

**Notes:**

1. If a file with an arrival sequence access path is reorganized using a keyed sequence access path, the arrival sequence access path is changed. That is, the records in the file are physically placed in the order of the keyed sequence access path used. By reorganizing the data into a physical sequence that closely matches the keyed access path you are using, you can improve the performance of processing the data sequentially.
2. Reorganizing a file compresses deleted records, which changes subsequent relative record numbers.
3. Because access paths with either the FIFO or LIFO DDS keyword specified depend on the physical sequence of records in the physical file, the sequence of the records with duplicate key fields may change after reorganizing a physical file using a keyed sequence access path.

Also, because access paths with the FCFO DDS keyword specified are ordered as FIFO, when a reorganization is done, the sequence of the records with duplicate key fields may also change.

4. If you cancel the RGZPFM command, all the access paths over the physical file member may have to be rebuilt.

If one of the following conditions occur and the Reorganize Physical File Member (RGZPFM) command is running, the records may not be reorganized:

- The system ends abnormally.
- The job containing the RGZPFM command is ended with an \*IMMED option.
- The subsystem in which the RGZPFM command is running ends with an \*IMMED option.
- The system stops with an \*IMMED option.

The status of the member being reorganized depends on how much the system was able to do before the reorganization was ended and what you specified in the SRCOPT parameter. If the SRCOPT parameter was not specified, the member is either completely reorganized or not reorganized at all. You should display the contents of the file, using the Display Physical File Member (DSPPFM) command, to determine if it was reorganized. If the member was not reorganized, issue the Reorganize Physical File Member (RGZPFM) command again.

If the SRCOPT parameter was specified, any of the following could have happened to the member:

- It was completely reorganized. A completion message is sent to your job log indicating the reorganize operation was completely successful.
- It was not reorganized at all. A message is sent to your job log indicating that the reorganize operation was not successful. If this occurs, issue the Reorganize Physical File Member (RGZPFM) command again.
- It was reorganized, but only some of the sequence numbers were changed. A completion message is sent to your job log indicating that the member was reorganized, but all the sequence numbers were not changed. If this occurs, issue the RGZPFM command again with KEYFILE(\*NONE) specified.

To reduce the number of deleted records that exist in a physical file member, the file can be created or changed to reuse deleted record space. For more information, see “Reusing Deleted Records” on page 5-3.

## Displaying Records in a Physical File Member

The Display Physical File Member (DSPPFM) command can be used to display the data in the physical database file members by arrival sequence. The command can be used for:

- Problem analysis
- Debugging
- Record inquiry

You can display source files or data files, regardless if they are keyed or arrival sequence. Records are displayed in arrival sequence, even if the file is a keyed file. You can page through the file, locate a particular record by record number, or shift the display to the right or left to see other parts of the records. You can also press a function key to show either character data or hexadecimal data on the display.

If you have Query installed, you can use the Start Query (STRQRY) command to select and display records, too.

If you have the SQL/400 language installed, you can use the Start SQL/400 (STRSQL) command to interactively select and display records.





---

## Chapter 11. Changing Database File Descriptions and Attributes

This chapter describes the things to consider when planning to change the description or attributes of a database file.

---

### Effect of Changing Fields in a File Description

When a program that uses externally described data is compiled, the compiler copies the file descriptions of the files into the compiled program. When you run the program, the system can verify that the record formats the program was compiled with are the same as the record formats currently defined for the file. The default is to do level checking.

The system assigns a unique level identifier for each record format when the file it is associated with is created. The system uses the information in the record format description to determine the level identifier. This information includes the total length of the record format, the record format name, the number and order of fields defined, the data type, the size of the fields, the field names, and the number of decimal positions in the field. Changes to this information in a record format cause the level identifier to change.

The following DDS information has no effect on the level identifier and, therefore, can be changed without recompiling the program that uses the file:

- TEXT keyword
- COLHDG keyword
- CHECK keyword
- EDTCDE keyword
- EDTWRD keyword
- REF keyword
- REFFLD keyword
- CMP, RANGE, and VALUES keywords
- TRNTBL keyword
- REFSHIFT keyword
- DFT keyword
- CCSID keyword
- ALWNULL keyword
- Join specifications and join keywords
- Key fields
- Access path keywords
- Select/omit fields

Keep in mind that even though changing key fields or select/omit fields will not cause a level check, the change may cause unexpected results in programs using the new access path. For example, changing the key field from the customer number to the customer name changes the order in which the records are retrieved, and may cause unexpected problems in the programs processing the file.

If level checking is specified (or defaulted to), the level identifier of the file to be used is compared to the level identifier of the file in your program when the file is opened. If the identifiers differ, a message is sent to the program to identify the

changed condition and the changes may affect your program. You can simply compile your program again so that the changes are included.

An alternative is to display the file description to determine if the changes affect your program. You can use the Display File Field Description (DSPFFD) command to display the description or, if you have SEU, you can display the source file containing the DDS for the file.

The format level identifier defined in the file can be displayed by the Display File Description (DSPFD) command. When you are displaying the level identifier, remember that the record format identifier is compared, rather than the file identifier.

Not every change in a file necessarily affects your program. For example, if you add a field to the end of a file and your program does not use the new field, you do not have to recompile your program. If the changes do not affect your program, you can use the Change Physical File (CHGPF) or the Change Logical File (CHGLF) commands with LVLCHK(\*NO) specified to turn off level checking for the file, or you can enter an Override with Database File (OVRDBF) command with LVLCHK(\*NO) specified so that you can run your program without level checking.

Keep in mind that level checking is the preferred method of operating. The use of LVLCHK(\*YES) is a good database integrity practice. The results produced by LVLCHK(\*NO) cannot always be predicted.

---

## Changing a Physical File Description and Attributes

Sometimes, when you make a change to a physical file description and then re-create the file, the level identifier can change. For example, the level identifier will change if you add a field to the file description, or change the length of an existing field. If the level identifier changes, you can compile the programs again that use the physical file. After the programs are recompiled, they will use the new level check identifier.

You can avoid compiling again by creating a logical file that presents data to your programs in the original record format of the physical file. Using this approach, the logical file has the same level check identifier as the physical file before the change.

For example, you decide to add a field to a physical file record format. You can avoid compiling your program again by doing the following:

1. Change the DDS and create a new physical file (FILEB in LIBA) to include the new field:

```
CRTPF FILE(LIBA/FILEB) MBR(*NONE)...
```

FILEB does not have a member. (The old file FILEA is in library LIBA and has one member MBRA.)

2. Copy the member of the old physical file to the new physical file:

```
CPYF FROMFILE(LIBA/FILEA) TOFILE(LIBA/FILEB)
FROMMBR(*ALL) TOMBR(*FROMMBR)
MBROPT(*ADD) FMOPT(*MAP)
```

The member in the new physical file is automatically named the same as the member in the old physical file because FROMMBR(\*ALL) and

TOMBR(\*FROMMBR) are specified. The FMTOPT parameter specifies to copy (\*MAP) the data in the fields by field name.

3. Describe a new logical file (FILEC) that looks like the original physical file (the logical file record format does *not* include the new physical file field). Specify FILEB for the PFILE keyword. (When a level check is done, the level identifier in the logical file and the level identifier in the program match because FILEA and FILEC have the same format.)

4. Create the new logical file:

```
CRTLF FILE(LIBA/FILEC)...
```

5. You can now do one of the following:

- a. Use an Override with Database File (OVRDBF) command in the appropriate jobs to override the old physical file referred to in the program with the logical file (the OVRDBF command parameters are described in more detail in Chapter 5, "Run Time Considerations").


```
OVRDBF FILE(FILEA) TOFILE(LIBA/FILEC)
```

- b. Delete the old physical file and rename the logical file to the name of the old physical file so the file name in the program does not have to be overridden.


```
DLTF FILE(LIBA/FILEA)
RNMOBJ OBJ(LIBA/FILEC) OBJTYPE(*FILE)
NEWOBJ(FILEA)
```

The following illustrates the relationship of the record formats used in the three files:

FILEA (old physical file)




FILEB (new physical file)



FLDB1 was added to the record format.

FILEC (logical file)



FILEC shares the record format of FILEA.

FLDB1 is not used in the record format for the logical file.

RSLH302-1

When you make changes to a physical file that cause you to create the file again, all logical files referring to it must first be deleted before you can delete and create the new physical file. After the physical file is re-created, you can re-create or restore the logical files referring to it. The following examples show two ways to do this:

**Example 1:** Create a new physical file with the same name in a different library

1. Create a new physical file with a different record format in a library different from the library the old physical file is in. The name of new file should be the same as the name of the old file. (The old physical file FILEPFC is in library LIBB and has two members, MBRC1 and MBRC2.)

```
CRTPF FILE(NEWLIB/FILEPFC) MAXMBS(2)...
```

2. Copy the members of the old physical file to the new physical file. The members in the new physical file are automatically named the same as the members in the old physical file because TOMBR(\*FROMMBR) and FROMMBR(\*ALL) are specified.

```
CPYF      FROMFILE(LIBB/FILEPFC) TOFILE(NEWLIB/FILEPFC)
          FROMMBR(*ALL) TOMBR(*FROMMBR)
          FMTOPT(*MAP *DROP) MBROPT(*ADD)
```

3. Describe and create a new logical file in a library different from the library the old logical file is in. The name of the new logical file should be the same as the old logical file name. You can use the FORMAT keyword to use the same record formats as in the current logical file if no changes need to be made to the record formats. You can also use the Create Duplicate Object (CRTDUPOBJ) command to create another logical file from the old logical file FILELFC in library LIBB.

```
CRTLFC FILE(NEWLIB/FILELFC)
```

4. Delete the old logical and physical files.

```
DLTF FILE(LIBB/FILELFC)
DLTF FILE(LIBB/FILEPFC)
```

5. Move the newly created files to the original library by using the following commands:

```
MOVOBJ OBJ(NEWLIB/FILELFC) OBJTYPE(*FILE) TOLIB(LIBB)
MOVOBJ OBJ(NEWLIB/FILEPFC) OBJTYPE(*FILE) TOLIB(LIBB)
```

**Example 2:** Creating new versions of files in the same libraries

1. Create a new physical file with a different record format in the same library the old physical file is in. The names of the files should be different. (The old physical file FILEPFA is in library LIBA and has two members MBRA1 and MBRA2.)

```
CRTPF FILE(LIBA/FILEPFB) MAXMBS(2)...
```

2. Copy the members of the old physical file to the new physical file.

```
CPYF      FROMFILE(LIBA/FILEPFA) TOFILE(LIBA/FILEPFB)
          FROMMBR(*ALL) TOMBR(*FROMMBR)
          FMTOPT(*MAP *DROP) MBROPT(*REPLACE)
```

3. Create a new logical file in the same library as the old logical file is in. The names of the old and new files should be different. (You can use the FORMAT keyword to use the same record formats as are in the current logical file if no changes need be made to the record formats.) The PFILE keyword must refer to the new physical file created in step 1. The old logical file FILELFA is in library LIBA.

```
CRTLFC FILE(LIBA/FILELFB)
```

4. Delete the old logical and physical files.

```
DLTF FILE(LIBA/FILELFA)
DLTF FILE(LIBA/FILEPFA)
```

5. Rename the new logical file to the name of the old logical file. (If you also decide to rename the physical file, be sure to change the DDS for logical file so that the PFILE keyword refers to the new physical file name.)

```
RNMOBJ(LIBA/FILELFB) OBJTYPE(*FILE) NEWOBJ(FILELFA)
```

6. If the logical file member should be renamed, and assuming the default was used on the Create Logical File (CRTLF) command, issue the following command:

```
RNMM FILE(LIBA/FILELFA) MBR(FILELFB) NEWMBR(FILELFA)
```

You can use the Change Physical File (CHGPF) command to change some of the attributes of a physical file and its members. For information on these parameters, see the Change Physical File (CHGPF) command in the *CL Reference* manual.

---

## Changing a Logical File Description and Attributes

As a general rule, when you make changes to a logical file that will cause a change to the level identifier (for example, adding a new field, deleting a field, or changing the length of a field), it is *strongly* recommended that you recompile the programs that use the logical file. Sometimes you can make changes to a file that change the level identifier and which do not require you to recompile your program (for example, adding a field that will not be used by your program to the end of the file). However, in those situations you will be forced to turn off level checking to run your program using the changed file. That is not the preferred method of operating. It increases the chances of incorrect data in the future.

To avoid recompiling, you can keep the current logical file (unchanged) and create a new logical file with the added field. Your program refers to the old file, which still exists.

You can use the Change Logical File (CHGLF) command to change most of the attributes of a logical file and its members that were specified on the Create Logical File (CRTLF) command.



---

## Chapter 12. Using Database Attribute and Cross-Reference Information

The AS/400 integrated database provides file attribute and cross-reference information. Some of the cross-reference information includes:

- The files used in a program
- The files that depend on other files for data or access paths
- File attributes
- The fields defined for a file

Each of the commands described in the following sections can present information on a display, a printout, or write the cross-reference information to a database file that, in turn, can be used by a program or utility (for example, Query) for analysis.

For more information about writing the output to a database file, see “Writing the Output from a Command Directly to a Database File” on page 12-4.

You can retrieve information about a member of a database file for use in your applications with the Retrieve Member Description (RTVMBRD) command. See the section on “Retrieving Member Description Information” in the *CL Programmer's Guide* for an example of how the RTVMBRD command is used in a CL program to retrieve the description of a specific member.

---

### Displaying Information about Database Files

#### Displaying Attributes for a File

You can use the Display File Description (DSPFD) command to display the file attributes for database files and device files. The information can be displayed, printed, or written to a database output file (OUTFILE). The information supplied by this command includes (parameter values given in parentheses):

- Basic attributes (\*BASATR)
- File attributes (\*ATR)
- Access path specifications (\*ACCPH, logical and physical files only)
- Select/omit specifications (\*SELECT, logical files only)
- Join logical file specifications (\*JOIN, join logical files only)
- Alternative collating sequence specifications (\*SEQ, physical and logical files only)
- Record format specifications (\*RCDFMT)
- Member attributes (\*MBR, physical and logical files only)
- Spooling attributes (\*SPOOL, printer and diskette files only)
- Member lists (\*MBRLIST, physical and logical files only)

#### Displaying the Descriptions of the Fields in a File

You can use the Display File Field Description (DSPFFD) command to display field information for both database and device files. The information can be displayed, printed, or written to a database output file (OUTFILE).

## Displaying the Relationships between Files on the System

You can use the Display Database Relations (DSPDBR) command to display the following information about the organization of your database:

- A list of database files (physical and logical) that use a specific record format.
- A list of database files (physical and logical) that depend on the specified file for data sharing.
- A list of members (physical and logical) that depend on the specified member for sharing data or sharing an access path.

This information can be displayed, printed, or written to a database output file (OUTFILE).

For example, to display a list of all database files associated with physical file ORDHDRP, with the record format ORDHDR, type the following DSPDBR command:

```
DSPDBR FILE(DSTPRODLB/ORDHDRP) RCDFMT(ORDHDR)
```

**Note:** See the DSPDBR command description in the *CL Reference* manual for details of this display.

This display presents header information when a record format name is specified on the RCDFMT parameter, and presents information about which files are using the specified record format.

If a member name is specified on the MBR parameter of the DSPDBR command, the dependent members are shown.

If the Display Database Relations (DSPDBR) command is specified with the default MBR(\*NONE) parameter value, the dependent data files are shown. To display the shared access paths, you must specify a member name.

The Display Database Relations (DSPDBR) command output identifies the type of sharing involved. If the results of the command are displayed, the name of the type of sharing is displayed. If the results of the command are written to a database file, the code for the type of sharing (shown below) is placed in the *WHTYPE* field in the records of the output file.

Type	Code	Description
Data	D	The file or member is dependent on the data in a member of another file.
Access path	I	(Access path sharing) The file member is sharing an access path.
Acc path owner	O	(Access path owner) If an access path is shared, one of the file members is considered the owner. The owner of the access path is charged with the storage used for the access path. If the member displayed is designated the owner, one or more file members are designated with an I for access path sharing.
SQL View	S	The SQL view or member is dependent upon another SQL view.



## Displaying the Files Used by Programs

You can use the Display Program Reference (DSPPGMREF) command to determine which files, data areas, and other programs are used by a program. This information is available for compiled programs only.

The information can be displayed, printed, or written to a database output file (OUTFILE).

When a program is created, the information about certain objects used in the program is stored. This information is then available for use with the Display Program References (DSPPGMREF) command.

The following chart shows the objects for which the high-level languages and utilities save information:

Language or Utility	Files	Programs	Data Areas	See Notes
BASIC	Yes	Yes	No	1
C/400 Language	No	No	N/A	
CL	Yes	Yes	Yes	2
COBOL/400 Language	Yes	Yes	No	3
CSP	Yes	Yes	No	4
DFU	Yes	N/A	N/A	
FORTRAN/400* Language	No	No	N/A	
Pascal	No	No	N/A	
PL/I	Yes	Yes	N/A	3
RPG/400 Language	Yes	Yes	Yes	5
SQL/400 Language	Yes	N/A	N/A	

### Notes:

1. Externally described file references, programs, and data areas are stored.
2. All system commands that refer to files, programs, or data areas specify in the command definition that the information should be stored when the command is compiled in a CL program. If a variable is used, the name of the variable is used as the object name (for example, &FILE). If an expression is used, the name of the object is stored as \*EXPR. User-defined commands can also store the information for files, programs, or data areas specified on the command. See the description of the FILE, PGM, and DTAARA parameters on the PARM or ELEM command statements in the *CL Programmer's Guide*.
3. The program name is stored only when a literal is used for the program name (this is a static call, for example, CALL 'PGM1'), not when a COBOL/400 identifier is used for the program name (this is a dynamic call, for example, CALL PGM1).
4. CSP programs also save information for an object of type \*MSGF, \*CSPMAP, and \*CSPTBL.
5. The use of the local data area is not stored.

The stored file information contains an entry (a number) for the type of use. In the database file output of the Display Program References (DSPPGMREF) command (built when using the OUTFILE parameter), this is specified as:

Code	Meaning
1	Input
2	Output
3	Input and Output
4	Update
8	Unspecified

Combinations of codes are also used. For example, a file coded as a 7 would be used for input, output, and update.

## Displaying the System Cross-Reference Files

The system manages two database files that contain basic file attribute information (QSYS/QADBXREF) and cross-reference information (QSYS/QADBFDEP) about all the database files on the system (except those database files that are in the QTEMP library). You can use these files to determine basic attribute and file requirements. To display the fields contained in these files, use the Display File Field Description (DSPFFD) command.

**Note:** Normally, the authority to use these files is restricted to the security officer.

---

## Writing the Output from a Command Directly to a Database File

You can store the output from many CL commands in an output physical file by specifying the OUTFILE parameter on the command.

You can use the output files in programs or utilities (for example, Query) for data analysis. For example, you can send the output of the Display Program References (DSPPGMREF) command to a physical file, then query that file to determine which programs use a specific file.

The physical files are created for you when you specify the OUTFILE parameter on the commands. Initially, the files are created with private authority; only the owner (the person who ran the command) can use it. However, the owner can authorize other users to these files as you would for any other database file.

The system supplies model files that identify the record format for each command that can specify the OUTFILE parameter. If you specify a file name on the OUTFILE parameter for a file that does not already exist, the system creates the file using the same record format as the model files. If you specify a file name for an existing output file, the system checks to see if the record format is the same record format as the model file. If the record formats do not match, the system sends a message to the job and the command does not complete.

**Note:** You must use your own files for output files, rather than specifying the system-supplied model files on the OUTFILE parameter.

See the *Programming Reference Summary* for a list of commands that allow output files and the names of the model files supplied for those commands.

**Note:** All system-supplied model files are located in the QSYS library.

You can display the fields contained in the record formats of the system-supplied model files using the Display File Field Descriptions (DSPFFD) command.

### Example of Using a Command Output File

The following example uses the Display Program References (DSPPGMREF) command to collect information for all compiled programs in all libraries, and place the output in a database file named DBROUT:

```
DSPPGMREF PGM(*ALL/*ALL) OUTPUT(*OUTFILE) OUTFILE(DSTPRODLB/DBROUT)
```

You can use Query to process the output file. Another way to process the output file is to create a logical file to select information from the file. The following is the DDS for such a logical file. Records are selected based on the file name.

```
|...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8
A* Logical file DBROUTL for query
A
A          R DBROUTL          PFILE(DBROUT)
A          S WHFNAM          VALUES('ORDHDRL' 'ORDFILL')
A
```

### Output File for the Display File Description Command

The Display File Description (DSPFD) command provides unique output files, depending on the parameters specified. See the *Programming Reference Summary* for a list of the model files for the DSPFD command.

**Note:** All system-supplied model files are in the QSYS library.

To collect access path information about all files in the LIBA library, you could specify:

```
DSPFD FILE(LIBA/*ALL) TYPE(*ACCPH) OUTPUT(*OUTFILE) +
      OUTFILE(LIBB/ABC)
```

The file ABC is created in library LIBB and is externally described with the same field descriptions as in the system-supplied file QSYS/QAFDACCP. The ABC file then contains a record for each key field in each file found in library LIBA that has an access path.

If the Display File Description (DSPFD) command is coded as:

```
DSPFD FILE(LIBX/*ALL) TYPE(*ATR) OUTPUT(*OUTFILE) +
      FILEATR(*PF) OUTFILE(LIBB/DEF)
```

the file DEF is created in library LIBB and is externally described with the same field descriptions as exist in QSYS/QAFDPHY. The DEF file then contains a record for each physical file found in library LIBX.

You can display the field names of each model file supplied by IBM using the DSPFFD command. For example, to display the field description for the access path model file (\*ACCPH specified on the TYPE parameter), specify the following:

```
DSPFFD QSYS/QAFDACCP
```

### Output Files for the Display Journal Command

See the *Programming Reference Summary* for a list of model output files supplied on the system that can be shown with the Display Journal (DSPJRN) command.

### Output Files for the Display Problem Command

See the *Programming Reference Summary* for a list of model output files supplied on the system for the Display Problem (DSPPRB) command. The command provides unique output files depending on the type of record:

- Basic problem data record (\*BASIC). This includes problem type, status, machine type/model/serial number, product ID, contact information, and tracking data.

- Point of failure, isolation, or answer FRU records (\*CAUSE). Answer FRUs are used if they are available. If answer FRUs are not available, isolation FRUs are used if they are available. If answer FRUs and isolation FRUs are not available, then point of failure FRUs are used.
- PTF fix records (\*FIX).
- User-entered text (note records) (\*USRTXT).
- Supporting data identifier records (\*SPTDTA).

The records in all five output files have a problem identifier so that the cause, fix, user text information, and supporting data can be correlated with the basic problem data. Only one type of data can be written to a particular output file. The cause, fix, user text, and supporting data output files can have multiple records for a particular problem. See the *CL Reference* manual for more information on the DSPPRB command.

---

## Chapter 13. Database Recovery Considerations

This chapter describes the general considerations and AS/400 facilities that enable you to recover or restore your database following any type of unexpected or undesirable event that could cause loss of data on the system. See the *Basic Backup and Recovery Guide* and the *Advanced Backup and Recovery Guide* for comprehensive discussion of AS/400 backup and recovery strategies, plans, and facilities.

---

### Database Save and Restore

It is important that you save your database files and related objects periodically so that you can restore them when necessary.

Database files and related objects can be saved and restored using diskette, magnetic tape, or a save file. When information is saved, a copy of the information in a special format is written onto one or more diskettes, reels of magnetic tape, tape cartridges, or to a save file. Diskettes and tape can be removed and stored for future use on your system or on another AS/400 system. When information is restored, it is read from diskette, tape, or a save file into storage where it can be accessed by system users.

Save files are disk-resident files that can be the target of a save operation or the source of a restore operation. Save files allow unattended save operations. That is, an operator does not need to load tapes or diskettes when saving to a save file. However, it is still important to use the Save Save File Data (SAVSAVFDTA) command to periodically save the save file data on tape or diskette. The tapes or diskettes should periodically be removed from the site. Storing a copy of your save tapes or diskettes away from the system site is important to help recover from a site disaster.

Saving your data can be done using dedicated save operations or a combination of dedicated and save-while-active operations. A **dedicated save** is a save operation that is scheduled when application programs are not running. **Save-while-active** allows saves to be done while application programs that update your database files and objects are running.

Refer to the *Basic Backup and Recovery Guide* for general planning and design considerations for save and restore and for information about using save and restore commands. Refer to the *Advanced Backup and Recovery Guide* for detail design and programming considerations related to using the save-while-active function.

### Considerations for Save and Restore

When you save a database file or related object to tape or diskette, the system updates the object description with the date and time of the save operation. When you save an object to a save file, you can prevent the system from updating the date and time of the save operation by specifying UPDHST(\*NO) on the save command. When you restore an object, the system always updates the object description with the date and time of the restore operation. You can display this and other save/restore related information by using the Display Object Description (DSPOBJD) command with DETAIL(\*FULL). Use the Display Save File (DSPSAVF) command to display the objects in a save file.

Specify DATA(\*SAVRST) on the Display Diskette (DSPDKT) or Display Tape (DSPTAP) command for a display of the objects on the media.

The last save/restore date for database members can be displayed by typing:

```
DSPFD FILE(file-name) TYPE(*MBR)
```

---

## Database Data Recovery

The AS/400 system has integrated recovery functions to help recover data in a database file. The key functions described in this chapter are:

- Journal management, for recording data changes to files
- Commitment control, for synchronizing transaction recovery
- Force-writing data changes to auxiliary storage
- Abnormal system end recovery (see “Database Recovery after an Abnormal System End” on page 13-9)

## Journal Management

Journal management allows you to record all the data changes occurring to one or more database files. You can then use the journal for recovery.

You should seriously consider using journal management. If a database file is destroyed or becomes unusable and you are using journaling, you can reconstruct most of the activity for the file (see journaling topic in the *Advanced Backup and Recovery Guide* for details). Optionally, the journal allows you to remove changes made to the file.

Journaling can be started or ended very easily. It requires no additional programming or changes to existing programs.

When a change is made to a file and you are using **journaling**, the system records the change in a journal receiver and writes the receiver to auxiliary storage before it is recorded in the file. Therefore, the journal receiver always has the latest database information. Activity for a file is journaled regardless of the type of program, user, or job that made the change, or the logical file through which the change was made.

Journal entries record activity for a specific record (record added, updated or deleted), or for the file as a whole (file opened, file member saved, and so on). Each entry includes additional bytes of control information identifying the source of the activity (including user, job, program, time, and date). For changes that affect a single record, record images are included following the control information. The image of the record after a change is made is always included. Optionally, the record image before the change is made can also be included. You control whether to journal both before and after record images or just after record images by specifying the IMAGES parameter on the Start Journaling Physical File (STRJRNPf) command.

All journaled database files are automatically synchronized with the journal when the system is started (IPL time). If the system ended abnormally, some database changes may be in the journal, but not yet reflected in the database itself. If that is the case, the system automatically updates the database from the journal to bring the database files up to date.

Journaling can make saving database files easier and faster. For example, instead of saving an entire file everyday, you can simply save the journal receiver that contains the changes to that file. You might still save the entire file on a weekly basis. This method can reduce the amount of time it takes to perform your daily save operations.

The Apply Journalled Changes (APYJRNCHG) and Remove Journalled Changes (RMVJRNCHG) commands can be used to recover a damaged or unusable database file member using the journalled changes. The APYJRNCHG command applies the changes that were recorded in a journal receiver to the designated physical file member. Depending on the type of damage to the physical file and the amount of activity since the file was last saved, removing changes from the file using the RMVJRNCHG command can be easier. The Work with Journal (WRKJRN) command provides a prompted method for applying and removing changes. See the *Advanced Backup and Recovery Guide* for descriptions and examples of these commands.

The Display Journal (DSPJRN) command, can be used to convert journal entries to a database file. Such a file can be used for activity reports, audit trails, security, and program debugging.

Because the journal supplies many useful functions, not the least of which is recovering data, journal management ought to be considered a key part of your recovery strategy. See the *Advanced Backup and Recovery Guide* for more information about journal management.

## Transaction Recovery through Commitment Control

**Commitment control** is a function that allows you to define and process a number of changes to database files as a single unit (transaction). Commitment control is an extension of the journal function on the system that provides you additional assistance in recovering data and restarting your programs. Commitment control can ensure that complex application transactions are logically synchronized even if the job or system ends.

A **transaction** is a group of changes that appear as a single change, such as the transfer of funds from a savings account to a checking account. Transactions can be classified as follows:

- Inquiries in which no file changes occur.
- Simple transactions in which one file is changed each time you press the Enter key.
- Complex transactions in which two or more files are changed each time you press the Enter key.
- Complex transactions in which one or more files are changed each time you press the Enter key, but these changes represent only part of a logical group of transactions.

Changes made to files during transaction processing are journalled when using commitment control.

If the system or job ends abnormally, journaling alone can ensure that, at most, only the very last record change is lost. However, if the system or job ends abnormally during a complex transaction (where more than one file may be changed), the files can reflect an incomplete logical transaction. For example, the job may have updated a record in file A, but before it had a chance to update a corresponding

record in file B the job ended abnormally. In this case, the logical transaction consisted of two updates, but only one update completed before the job ended abnormally.

Recovering a complex application requires detailed application knowledge. Programs cannot simply be restarted; for example, record changes may have to be made with an application program or data file utility to reverse the files to just before the last complex transaction began. This task becomes more complex if multiple users were accessing the files at the same time.

Commitment control helps solve these problems. Under commitment control, the records used during a complex transaction are locked from other users. This ensures that other users do not use the records until the transaction is complete. At the end of the transaction, the program issues the commit operation, freeing the records. However, should the system or job end abnormally before the commit operation is performed, all record changes for that job since the last time a commit operation occurred are rolled back. Any affected records that are still locked are then unlocked. In other words, database changes are rolled back to a clean transaction boundary.

The rollback operation can also occur under your control. Assume that in an order entry application, the application program runs the commit operation at the end of each order. In the middle of an order, the operator can signal the program to do a rollback operation. All file changes will be rolled back to the beginning of the order.

The commit and roll back operations are available in several AS/400 programming languages including the RPG/400, COBOL/400, PL/I, SQL/400, and the AS/400 control language (CL).

An optional feature of commitment control is the use of a notify object. The notify object is a file, data area, or message queue. When the job or system ends during a transaction, information specified by the program is automatically sent to the notify object. This information can be used by an operator or application programs to start the application from the last successful transaction boundary.

Commitment control can also be used in a batch environment. Just as it provides assistance in interactive transaction recovery, commitment control can help in batch job recovery. See the *Advanced Backup and Recovery Guide* for more information about commitment control.

## Force-Writing Data to Auxiliary Storage

The force-write ratio (FRCRATIO) parameter on the create file and override database file commands can be used to force data to be physically written to auxiliary storage. A force-write ratio of one causes every add, update, and delete request to be immediately written to auxiliary storage for the file in question. However, choosing this option can reduce system performance. Therefore, saving your files and journaling your files should be considered the primary methods for protecting database files.



---

## Access Path Recovery

The system ensures the integrity of an access path before you can use it. If the system determines that the access path is unusable, the system attempts to recover it. You can control when an access path will be recovered. See “Controlling When Access Paths Are Rebuilt” on page 13-6 for more information.

Access path recovery can take a long time, especially if you have large access paths or many access paths to be rebuilt. You can reduce this recovery time in several ways, including:

- Saving access paths
- Using access path journaling
- Controlling when access paths are rebuilt
- Designing files to reduce rebuild time

## Saving Access Paths

You can reduce the time required to recover access paths by saving access paths. The access path (ACCPATH) parameter on the SAVCHGOBJ, SAVLIB, and SAVOBJ commands allows you to save access paths. Normally, only the descriptions of logical files are saved; however, the access paths are saved under the following conditions:

- ACCPTH(\*YES) is specified.
- All physical files under the logical file are being saved and are in the same library.
- The logical file is MAINT(\*IMMED) or MAINT(\*DLY).

See the *Basic Backup and Recovery Guide* for additional information.

## Restoring Access Paths

Access paths can be restored if they were saved and if all the physical files on which they depend are restored at the same time. See the *Basic Backup and Recovery Guide* for additional information.

Restoring an access path can be faster than rebuilding it. For example, assume a logical file is built over a physical file containing 500,000 records and you have determined (through the Display Object Description [DSPOBJD] command) that the size of the logical file is about 15 megabytes. In this example, assume it would take about 50 minutes to rebuild the access path for the logical file compared to about 1 minute to restore the same access path from a tape. (This assumes that the system can build approximately 10,000 index entries per minute.)

After restoring the access path, the file may need to be brought up-to-date by applying the latest journal changes (depending on whether journaling is active). Normally, the system can apply approximately 80,000 to 100,000 journal entries per hour. (This assumes that each of the physical files to which entries are being applied has only one access path built over it.) Even with this additional recovery time, you will usually find it is faster to restore access paths rather than to rebuild them.

## Rebuilding Access Paths

Rebuilding a database access path may take as much as one minute for every 10,000 records.

**Note:** This estimate should be used until actual times for your system can be calculated.

The following factors affect this time estimate (listed in general order of significance):

- Storage pool size. The size of the storage pool used to rebuild the access path is a very important factor. You can improve the rebuild time by running the job in a larger storage pool.
- The system model. The speed of the processing unit is a key factor in the time needed to rebuild an access path.
- Key length. A large key length will slow rebuilding the access path because more key information must be constructed and stored in the access path.
- Select/omit values. Select/omit processing will slow the rebuilding of an access path because each record must be compared to see if it meets the select/omit values.
- Record length. A large record length will slow the rebuilding of an access path because more data is looked at.
- Storage device containing the data. The relative speed of the storage device containing the actual data and the device where the access path is stored has an effect on the time needed to rebuild an access path.
- The order of the records in the file. The system tries to rebuild an access path so that it can find information quickly when using that access path. The order of the records in a file has a small affect on how fast the system can build the access path while trying to maintain an efficient access path.

All of the preceding factors must be considered when estimating the amount of time to rebuild an access path.

## Controlling When Access Paths Are Rebuilt

If the system ends abnormally, during the next IPL the system automatically lists those files requiring access path recovery. You can decide whether to rebuild the access path:

- During the IPL
- After the IPL
- When the file is first used

You can also:

- Change the scheduling order in which the access paths are rebuilt
- Hold rebuilding of an access path indefinitely
- Continue the IPL process while access paths with a sequence value that is less than or equal to the *\*IPL threshold* value are rebuilding
- Control the rebuilding of access paths after the system has completed the IPL process by using the Edit Rebuild of Access Paths (EDTRBDAP) command

The IPL threshold value is used to determine which access paths rebuild during the IPL. All access paths with a sequence value that is less than or equal to the IPL threshold value rebuild during the IPL. Changing the IPL threshold value to 99 means that all access paths with a sequence value of 1 through 99 rebuild during the IPL. Changing the IPL threshold value to 0 means that no access paths rebuild until after the system completes its IPL, except access paths that were being journaled and access paths for system files.

The access path recovery value for a file is determined by the value you specified for the RECOVER parameter on the create and change file commands. The default recovery value for \*IPL (rebuild during IPL) is 25 and the default value for \*AFTIPL (rebuild after IPL) is 75; therefore, RECOVER(\*IPL) will show as 25. The initial IPL threshold value is 50; this allows the parameters to affect when the access path is rebuilt. You can override this value on the Edit Rebuild of Access Paths display.

If a file is not needed immediately after IPL time, specify that the file can be rebuilt at a later time. This should help reduce the number of files that need to be rebuilt at IPL, allowing the system to complete its IPL much faster.

For example, you can specify that all files that must have their access paths rebuilt should rebuild the access paths when the file is first used. In this case, no access paths are rebuilt at IPL. You can control the order in which the access paths are rebuilt by running only those programs that use the files you want rebuilt first. This method shortens the IPL time (because there are no access paths to rebuild during the IPL) and could make the first of several applications available faster. However, the overall time to rebuild all the access paths probably is longer (because there may be other work running when the access paths are being rebuilt, and there may be less main storage available to rebuild the access paths).

### **Designing Files to Reduce Access Path Rebuilding Time**

File design can also help reduce access path recovery time. For example, you might divide a large master file into a history file and a transaction file. The transaction file would be used for adding new data, the history file would be used for inquiry only. On a daily basis, you might merge the transaction data into the history file, then clear the transaction file for the next day's data. With this design, the time to rebuild access paths could be shortened. That is, if the system abnormally ended during the day the access path to the smaller transaction file might need to be rebuilt. However, the access path to the large history file, being read-only for most of the day, would rarely be out of synchronization with its data, thereby significantly reducing the chances that it would have to be rebuilt.

Consider the trade-off between using a file design to reduce access path rebuilding time and using system-supplied functions like access path journaling. The file design described above may require a more complex application design. After evaluating your situation, you may decide to use system-supplied functions like access path journaling rather than design more complex applications.

## Journaling Access Paths

Journaling access paths can significantly reduce recovery time by reducing the number of access paths that need to be rebuilt after an abnormal system end.

**Note:** Journaling access paths is strongly recommended for AS/400 Version 2 Release 2 and following releases, because access paths may become much larger and may therefore require more time to rebuild.

When you journal database files, images of changes to the records in the file are recorded in the journal. These record images are used to recover the file should the system end abnormally. However, after an abnormal end, the system may find that access paths built over the file are not synchronized with the data in the file. If an access path and its data are not synchronized, the system must rebuild the access path to ensure that the two are synchronized and usable.

When access paths are journaled, the system records images of the access path in the journal to provide known synchronization points between the access path and its data. By having that information in the journal, the system can recover both the data files and the access paths, and ensure that the two are synchronized. In such cases, the lengthy time to rebuild the access paths can be avoided.

In addition, journaling access paths works with other recovery functions on the system. For example, the system has a number of options to help reduce the time required to recover from the failure and replacement of a disk unit. These options include user auxiliary storage pools and checksum protection. While these options reduce the chances that the entire system must be reloaded because of the disk failure, they do not change the fact that access paths may need to be rebuilt when the system is started following replacement of the failed disk. By using access path journaling and some of the recovery options discussed previously, you can reduce your chances of having to reload the entire system and having to rebuild access paths.

Journaling access paths is easy to start. The Start Journal Access Path (STRJRNAP) command is used to start journaling the access path for the specified file. You can journal access paths that have a maintenance attribute of immediate (\*IMMED) or delayed (\*DLY). Once journaling is started, the system continues to protect the access path until the access path is deleted or you run the End Journaling Access Path (ENDJRNAP) command for that access path.

Before journaling an access path, you must journal the physical files associated with the access path. In addition, you must use the same journal for the access path and its associated physical files.

Access path journaling is designed to minimize additional output operations. For example, the system will write the journal data for the changed record and the changed access path in the same output operation. However, you should seriously consider isolating your journal receivers in user auxiliary storage pools when you start journaling your access paths. Placing journal receivers in their own user auxiliary storage pool provides the best journaling performance, while helping to protect them from a disk failure. See the *Advanced Backup and Recovery Guide* for more information about journaling access paths.

## Other Methods to Avoid Rebuilding Access Paths

If you do not journal your access paths, then you might consider some other system functions that can help you reduce the chances of having to rebuild access paths.

The method used by the system to determine if an access path needs to be rebuilt is a file synchronization indicator. Normally the synchronization indicator is on, indicating that the access path and its associated data are synchronized. When a job changes a file that affects an access path, the system turns off the synchronization indicator in the file. If the system ends abnormally, it must rebuild any access path whose file has its synchronization indicator off.

To reduce the number of access paths that must be rebuilt, you need a way to periodically synchronize the data with its access path. There are several methods to synchronize a file with its access path:

- Full file close. The last full (that is, not shared) system-wide close performed against a file will synchronize the access path and the data.
- Force access path. The force-access-path (FRCACCPH) parameter can be specified on the create or change file commands.
- Force write ratio of 2 or greater. The force-write-ratio (FRCRATIO) parameter can be specified on the create, change, or override database file commands.
- Force end of data. The file's data and its access path can be synchronized by running the force-end-of-data operation in your program. (Some high-level languages do not have a force-end-of-data operation. See your high-level language guide for further details.)

Keep in mind that while the data and its access path are synchronized after performing one of the methods mentioned previously, the next change to the data in the file can cause the synchronization indicator to be turned off again. It is also important to note that each of the methods can be costly in terms of performance; therefore, they should be used with caution. Consider journaling access paths, along with saving access paths, as the primary means of protecting access paths.

---

## Database Recovery after an Abnormal System End

After an abnormal system end, the system proceeds through several automatic recovery steps. This includes such things as: rebuilding the system directory and synchronizing the journal to the files being journaled. The system performs recovery operations during IPL and after IPL.

### Database File Recovery during the IPL

During IPL, nothing but the recovery function is active on the system. During IPL, database file recovery consists of the following:

- The following functions that were in progress when the system ended are completed:
  - Delete file
  - Remove member
  - Rename member
  - Move object
  - Rename object

- Change object owner
  - Change file
  - Change member
  - Grant authority
  - Revoke authority
  - Start journaling physical file
  - Start journaling access path
  - End journaling physical file
  - End journaling access path
  - Change journal
  - Delete journal
  - Recover SQL/400 views
- The following functions that were in progress when the system ended are backed out (you must run them again):
    - Create file
    - Add member
    - Create journal
    - Restore journal
  - If the operator is doing the IPL (attended IPL), the Edit Rebuild of Access Paths display appears on the operator's display. The display allows the operator to edit the RECOVER option for the files that were in use for which immediate or delayed maintenance was specified. (See "Controlling When Access Paths Are Rebuilt" on page 13-6 for information about the IPL threshold value.) If all access paths are valid, or the IPL is unattended, no displays appear.
  - Access paths that have immediate or delayed maintenance, and that are specified for recovery during IPL (from the RECOVER option or changed by the Edit Rebuild of Access Paths display) are rebuilt and a message is sent to the history log. Files with journaled access paths that were in use, and system files with access paths that are not valid, are not displayed on the Edit Rebuild of Access Paths display. They are automatically recovered and a message is sent to the history log.
  - For unattended IPLs, if the system value QDBRCVYWT is 1 (wait), files that were in use that are specified for recovery after IPL are treated as files specified for recovery during IPL. See the *Work Management Guide* for more information on the system value QDBRCVYWT.
  - Messages about the following information are sent to the history log:
    - The success or failure of the previously listed items
    - The physical file members that were open when the system ended abnormally and the last active relative record number in each member
    - The physical file members that could not be synchronized with the journal
    - That IPL database recovery has completed

### **Database File Recovery after the IPL**

This recovery step is run after the IPL is complete. Interactive users may be active and batch jobs may be running with this step of database recovery.

Recovery after the IPL consists of the following:

- The access paths for immediate or delayed maintenance files which specify recovery after IPL, are rebuilt (see Figure 13-1 on page 13-11).

- Messages are sent to the system history log indicating the success or failure of the rebuild operation.
- After the IPL completes, the Edit Rebuild of Access Paths (EDTRBDAP) command can be used to order the rebuilding of access paths.

**Note:** If you are not using journaling for a file, records may or may not exist after IPL recovery, as follows:

- For added records, if after the IPL recovery the Nth record added exists, then all records added preceding N also exist.
- For updated and deleted records, if the update or delete to the Nth record is present after the IPL recovery, there is no guarantee that the records updated or deleted prior to the Nth record are also present in the database.
- For REUSEDLT(\*YES), records added are treated as updates, and thus, there is no guarantee that records exist after IPL recovery.

### Database File Recovery Options Table

The table below summarizes the file recovery options:

Figure 13-1. Relationship of Access Path, Maintenance, and Recovery

Access Path/ Maintenance	RECOVER Parameter Specified		
	*NO	*AFTIPL	*IPL
Keyed sequence access path/ immediate or delayed maintenance	<ul style="list-style-type: none"> <li>• No database recovery at IPL</li> <li>• File available immediately</li> <li>• Access path rebuilt first time file opened</li> </ul>	<ul style="list-style-type: none"> <li>• Access path rebuilt after IPL</li> </ul>	<ul style="list-style-type: none"> <li>• Access path rebuilt during IPL</li> </ul>
Keyed sequence access path rebuild maintenance	<ul style="list-style-type: none"> <li>• No database recovery at IPL</li> <li>• File available immediately</li> <li>• Access path rebuilt first time file opened</li> </ul>	<ul style="list-style-type: none"> <li>• Not applicable; no recovery is done for rebuild maintenance</li> </ul>	<ul style="list-style-type: none"> <li>• Not applicable; no recovery is done for rebuild maintenance</li> </ul>
Arrival sequence access path	<ul style="list-style-type: none"> <li>• No database recovery at IPL</li> <li>• File available immediately</li> </ul>	<ul style="list-style-type: none"> <li>• Not applicable; no recovery is done for an arrival sequence access path</li> </ul>	<ul style="list-style-type: none"> <li>• Not applicable; no recovery is done for an arrival sequence access path</li> </ul>

### Storage Pool Paging Option Effect on Database Recovery

The shared pool paging option controls whether the system dynamically adjusts the paging characteristics of the storage pool for optimum performance.

- The system does not dynamically adjust paging characteristics for a paging option of \*FIXED.
- The system dynamically adjusts paging characteristics for a paging option of \*CALC.
- You can also control the paging characteristics through an application programming interface. For more information, see Change Pool Tuning Information

| API(QWCCHGTN) in the *System Programmer's Interface Reference*,  
| SC41-8223

| A shared pool paging option other than \*FIXED can have an impact on data loss  
| for nonjournaled physical files in a system failure. When you do not journal phys-  
| ical files, data loss from a system failure, where memory is not saved, can increase  
| for \*CALC or USRDFN paging options. File changes may be written to auxiliary  
| storage less frequently for these options. There is a risk of data loss for non-  
| journaled files with the \*FIXED option, but the risk can be higher for \*CALC or user  
| defined (USRDFN) paging options.

| For more information on the paging option see the "Automatic System Tuning"  
| section of the *Work Management Guide*



---

## Chapter 14. Using Source Files

This chapter describes source files. Source file concepts are discussed, along with why you would use a source file. In addition, this chapter describes how to set up a source file, how to enter data into a source file, and how to use that source file to create another object (for example, a file or program) on the system.

---

### Source File Concepts

A source file is used when a command alone cannot supply sufficient information for creating an object. It contains input (source) data needed to create some types of objects. For example, to create a control language (CL) program, you must use a source file containing source statements, which are in the form of commands. To create a logical file, you must use a source file containing DDS.

To create the following objects, source files are required:

- High-level language programs
- Control language programs
- Logical files
- Intersystem communications function (ICF) files
- Commands
- Translate tables

To create the following objects, source files can be used, but are *not* required:

- Physical files
- Display files
- Printer files

A source file can be a database file, diskette file, tape file, or inline data file. (An inline data file is included as part of a job.) A source database file is simply another type of database file. You can use a source database file like you would any other database file on the system.

---

### Creating a Source File

To create a source file, you can use the Create Source Physical File (CRTSRCPF), Create Physical File (CRTPF), or Create Logical File (CRTLFL) command. Normally, you will use the CRTSRCPF command to create a source file, because many of the parameters default to values that you usually want for a source file. (If you want to use DDS to create a source file, then you would use the CRTPF or CRTLFL command.)

The CRTSRCPF command creates a physical file, but with attributes appropriate for source physical files. For example, the default record length for a source file is 92 (80 for the source data field, 6 for the source sequence number field, and 6 for the source date field).

The following example shows how to create a source file using the CRTSRCPF command and using the command defaults:

```
CRTSRCPF FILE(QGPL/FRSOURCE) TEXT('Source file')
```

## IBM-Supplied Source Files

For your convenience, the OS/400 program and other licensed programs provide a database source file for each type of source. These source files are:

File Name	Library Name	Used to Create
QBASSRC	QGPL	BASIC programs
QCBLSRC	QGPL	System/38 compatible COBOL
QCSRC	QGPL	C programs
QCLSRC	QGPL	CL programs
QCMDSRC	QGPL	Command definition statements
QFTNSRC	QGPL	FORTTRAN programs
QDDSSRC	QGPL	Files
QFMTSRC	QGPL	Sort source
QLBLSRC	QGPL	COBOL/400 programs
QS36SRC	#LIBRARY	System/36 compatible COBOL programs
QAPLISRC	QPLI	PL/I programs
QPLISRC	QGPL	PL/I programs
QREXSRC	QGPL	Procedures Language 400/REXX programs
QRPGSRC	QRPG	RPG/400 programs
QARPGSRC	QRPG38	System/38 environment RPG
QRPG2SRC	#RPGLIB	System/36 compatible RPG II
QS36SRC	#LIBRARY	System/36 compatible RPG II (after install)
QPASSRC	QPAS	Pascal programs
QTBLSRC	QGPL	Translation tables
QTXTSRC	QPDA	Text

You can either add your source members to these files or create your own source files. Normally, you will want to create your own source files using the same names as the IBM-supplied files, but in different libraries (IBM-supplied files may get overlaid when a new release of the system is installed). The IBM-supplied source files are created with the file names used for the corresponding create command (for example, the CRTCLPGM command uses the QCLSRC file name as the default). Additionally, the IBM-supplied programmer menu uses the same default names. If you create your own source files, do not place them in the same library as the IBM-supplied source files. (If you use the same file names as the IBM-supplied names, you should ensure that the library containing your source files precedes the library containing the IBM-supplied source files in the library list.)

## Source File Attributes

Source files usually have the following attributes:

- A record length of 92 characters (this includes a 6-byte sequence number, a 6-byte date, and 80 bytes of source).
- Keys (sequence numbers) that are unique even though the access path does not specify unique keys. You are not required to specify a key for a source file. Default source files are created without keys (arrival sequence access path). A source file created with an arrival sequence access path requires less storage space and reduces save/restore time in comparison to a source file for which a keyed sequence access path is specified.
- More than one member.
- Member names that are the same as the names of the objects that are created using them.
- The same record format for all records.
- Relatively few records in each member compared to most data files.

Some restrictions are:

- The source sequence number must be used as a key, if a key is specified.
- The key, if one is specified, must be in ascending sequence.
- The access path cannot specify unique keys.
- The ALTSEQ keyword is not allowed in DDS for source files.
- The first field must be a 6-digit sequence number field containing zoned decimal data and two decimal digits.
- The second field must be a 6-digit date field containing zoned decimal data and zero decimal digits.
- All fields following the second field must be zoned decimal or character.

## Creating Source Files without DDS

When you create a source physical file without using DDS, but by specifying the record length (RCDLEN parameter), the source created contains three fields: *SRCSEQ*, *SRCDAT*, and *SRCDTA*. (The record length must include 12 characters for sequence number and date-of-last-change fields so that the length of the data portion of the record equals the record length minus 12.) The data portion of the record can be defined to contain more than one field (each of which must be character or zoned decimal). If you want to define the data portion of the record as containing more than one field, you must define the fields using DDS.

A record format consisting of the following three fields is automatically used for a source physical file created using the Create Source Physical File (CRTSRCPF) command:

Field	Name	Data Type and Length	Description
1	SRCSEQ	Zoned decimal, 6 digits, 2 decimal positions	Sequence number for record
2	SRCDAT	Zoned decimal, 6 digits, no decimal positions	Date of last update of record
3	SRCDTA	Character, any length	Data portion of the record (text)

**Note:** For all IBM-supplied database source files, the length of the data portion is 80 bytes. For IBM-supplied device source files, the length of the data portion is the maximum record length for the associated device.

## Creating Source Files with DDS

If you want to create a source file for which you need to define the record format, use the Create Physical File (CRTPF) or Create Logical File (CRTLFL) command. If you create a source logical file, the logical file member should only refer to one physical file member to avoid duplicate keys.

---

## Working with Source Files

The following section describes how to enter and maintain data in your source files.

## Using the Source Entry Utility

You can use SEU, to enter and change source in a source file. If you use SEU to enter source in a database file, SEU adds the sequence number and date fields to each source record. (See the *SEU User's Guide and Reference* for more information about SEU.)

If you use SEU to update a source file, you can add records between existing records. For example, if you add a record between records 0003.00 and 0004.00, the sequence number of the added record could be 0003.01. SEU will automatically arrange the newly added statements in this way.

When records are first placed in a source file, the date field is all zoned decimal zeros (unless DDS is used with the DFT keyword specified). If you use SEU, the date field changes in a record when you change the record.

## Using Device Source Files

Tape and diskette unit files can be created as source files. When device files are used as source files, the record length must include the sequence number and date fields. Any maximum record length restrictions must consider these additional 12 characters. For example, the maximum record length for a tape record is 32 766. If data is to be processed as source input, the actual tape data record has a maximum length of 32 754 (which is 32 766 minus 12).

If you open source device files for input, the system adds the sequence number and date fields, but there are zeros in the date fields.

If you open a device file for output and the file is defined as a source file, the system deletes the sequence number and date before writing the data to the device.

## Copying Source File Data

The Copy Source File (CPYSRCF) and Copy File (CPYF) commands can be used to write data to and from source file members.

### Using the Copy Source File (CPYSRCF) Command for Copying to and from Source Files

The CPYSRCF command is designed to operate with database source files. Although it is similar in function to the Copy File (CPYF) command, the CPYSRCF command provides defaults that are normally used when copying a source file. For example, it has a default that assumes the TOMBR parameter is the same as the FROMMBR parameter and that any TOMBR records will always be replaced. The CPYSRCF command also supports a unique printing format when TOFILE(\*PRINT) is specified. Therefore, when you are copying database source files, you will probably want to use the CPYSRCF command.

The CPYSRCF command automatically converts the data from the from-file CCSID to the to-file CCSID.

## Using the Copy File (CPYF) Command for Copying to and from Files

The CPYF command provides additional functions over the CPYSRCF command such as:

- Copying from database source files to device files
- Copying from device files to database source files
- Copying between database files that are not source files and source database files
- Printing a source member in hexadecimal format
- Copying source with selection values

### Source Sequence Numbers Used in Copies

When you copy to a database source file, you can use the SRCOPT parameter to update sequence numbers and initialize dates to zeros. By default, the system assigns a sequence number of 1.00 to the first record and increases the sequence numbers by 1.00 for the remaining records. You can use the SRCSEQ parameter to set a fractional increased value and to specify the sequence number at which the renumbering is to start. For example, if you specify in the SRCSEQ parameter that the increased value is .10 and is to start at sequence number 100.00, the copied records have the sequence numbers 100.00, 100.10, 100.20, and so on.

If a starting value of .01 and an increased value of .01 are specified, the maximum number of records that can have unique sequence numbers is 999,999. When the maximum sequence number (9999.99) is reached, any remaining records will have a sequence number of 9999.99.

The following is an example of copying source from one member to another in the same file. If MBRB does not exist, it is added; if it does exist, all records are replaced.

```
CPYSRCF FROMFILE(QCLSRC) TOFILE(QCLSRC) FROMMBR(MBRA) +  
        TOMBR(MBRB)
```

The following is an example of copying a generic member name from one file to another. All members starting with PAY are copied. If the corresponding members do not exist, they are added; if they do exist, all records are replaced.

```
CPYSRCF FROMFILE(LIB1/QCLSRC) TOFILE(LIB2/QCLSRC) +  
        FROMMBR(PAY*)
```

The following is an example of copying the member PAY1 to the printer file QSYSPRT (the default for \*PRINT). A format similar to the one used by SEU is used to print the source statements.

```
CPYSRCF FROMFILE(QCLSRC) TOFILE(*PRINT) FROMMBR(PAY1)
```

When you copy from a device source file to a database source file, sequence numbers are added and dates are initialized to zeros. Sequence numbers start at 1.00 and are increased by 1.00. If the file being copied has more than 9999 records, then the sequence number is wrapped back to 1.00 and continues to be increased unless the SRCOPT and SRCSEQ parameters are specified.

When you are copying from a database source file to a device source file, the date and sequence number fields are removed.

## Using Source Files in a Program

You can process a source file in your program. You can use the external definition of the source file and do any input/output operations for that file, just as you would for any other database file.

Source files are externally described database files. As such, when you name a source file in your program and compile it, the source file description is automatically included in your program printout. For example, assume you wanted to read and update records for a member called FILEA in the source file QDDSSRC. When you write the program to process this file, the system will include the *SRCSEQ*, *SRCDAT*, and *SRCDTA* fields from the source file.

**Note:** You can display the fields defined in a file by using the Display File Field Description command (DSPFFD). For more information about this command, see “Displaying the Descriptions of the Fields in a File” on page 12-1.

The program processing the FILEA member of the QDDSSRC file could:

- Open the file member (just like any other database file member).
- Read and update records from the source file (probably changing the *SRCDTA* field where the actual source data is stored).
- Close the source file member (just like any other database file member).

---

## Creating an Object Using a Source File

You can use a create command to create an object using a source file. If you create an object using a source file, you can specify the name of the source file on the create command.

For example, to create a CL program, you use the Create Control Language Program (CRTCLPGM) command. A create command specifies through a SRCFILE parameter where the source is stored.

The create commands are designed so that you do not have to specify source file name and member name if you do the following:

1. Use the default source file name for the type of object you are creating. (To find the default source file name for the command you are using, see “IBM-Supplied Source Files” on page 14-2.)
2. Give the source member the same name as the object to be created.

For example, to create the CL program PGMA using the command defaults, you would simply type:

```
CRTCLPGM PGM(PGMA)
```

The system would expect the source for PGMA to be in the PGMA member in the QCLSRC source file. The library containing the QCLSRC file would be determined by the library list.

As another example, the following Create Physical File (CRTPF) command creates the file DSTREF using the database source file FRSOURCE. The source member is named DSTREF. Because the SRCMBR parameter is not specified, the system assumes that the member name, DSTREF, is the same as the name of the object being created.

CRTPF FILE (QGPL/DSTREF) SRCFILE(QGPL/FRSOURCE)

## Creating an Object from Source Statements in a Batch Job

If your create command is contained in a batch job, you can use an inline data file as the source file for the command. However, inline data files used as a source file should not exceed 10,000 records. The inline data file can be either named or unnamed. Named inline data files have a unique file name that is specified on the //DATA command. (For more information about inline data files, see the *Data Management Guide*.)

Unnamed inline data files are files without unique file names; they are all named QINLINE. The following is an example of an inline data file used as a source file:

```
//BCHJOB
CRTPF FILE(DSTPRODLB/ORD199) SRCFILE(QINLINE)
//DATA FILETYPE(*SRC)
.
.   (source statements)
.
//
//ENDBCHJOB
```

In this example, no file name was specified on the //DATA command. An unnamed spooled file was created when the job was processed by the spooling reader. The CRTPF command must specify QINLINE as the source file name to access the unnamed file. The //DATA command also specifies that the inline file is a source file (\*SRC specified for the FILETYPE parameter).

If you specify a file name on the //DATA command, you must specify the same name on the SRCFILE parameter on the CRTPF command. For example:

```
//BCHJOB
CRTPF FILE(DSTPRODLB/ORD199) SRCFILE(ORD199)
//DATA FILE(ORD199) FILETYPE(*SRC)
.
.   (source statements)
.
//
//ENDBCHJOB
```

If a program uses an inline file, the system searches for the first inline file of the specified name. If that file cannot be found, the program uses the first file that is unnamed (QINLINE).

If you do not specify a source file name on a create command, an IBM-supplied source file is assumed to contain the needed source data. For example, if you are creating a CL program but you did not specify a source file name, the IBM-supplied source file QCLSRC is used. You must have placed the source data in QCLSRC.

If a source file is a database file, you can specify a source member that contains the needed source data. If you do not specify a source member, the source data must be in a member that has the same name as the object being created.

## Determining Which Source File Member Was Used to Create an Object

When an object is created from source, the information about the source file, library, and member is held in the object. The date/time that the source member was last changed before object creation is also saved in the object.

The information in the object can be displayed with the Display Object Description (DSPOBJD) command and specifying DETAIL(\*SERVICE).

This information can help you in determining which source member was used and if the existing source member was changed since the object was created.

You can also ensure that the source used to create an object is the same as the source that is currently in the source member using the following commands:

- The Display File Description (DSPFD) command using TYPE(\*MBR). This display shows both date/times for the source member. The Last source update date/time value should be used to compare to the Source file date/time value displayed from the DSPOBJD command.
- The Display Object Description (DSPOBJD) command using DETAIL(\*SERVICE). This display shows the date/time of the source member used to create the object.

**Note:** If you are using the data written to output files to determine if the source and object dates are the same, then you can compare the *ODSRCD* (source date) and *ODSRCT* (source time) fields from the output file of the DSPOBJD DETAIL(\*SERVICE) command to the *MBUPDD* (member update date) and *MBUPDT* (member update time) fields from the output file of the DSPFD TYPE(\*MBR) command.

---

## Managing a Source File

This section describes several considerations for managing source files.

### Changing Source File Attributes

If you are using SEU to maintain database source files, see the *SEU User's Guide and Reference* for how to change database source files. If you are not using SEU to maintain database source files, you must totally replace the existing member.

If your source file is on a diskette, you can copy it to a database file, change it using SEU, and copy it back to a diskette. If you do not use SEU, you have to delete the old source file and create a new source file.

If you change a source file, the object previously created from the source file does not match the current source. The old object must be deleted and then created again using the changed source file. For example, if you change the source file FRSOURCE created in "Creating an Object Using a Source File" on page 14-6, you have to delete the file DSTREF that was created from the original source file, and create it again using the new source file so that DSTREF matches the changed FRSOURCE source file.



## Reorganizing Source File Member Data

You usually do not need to reorganize a source file member if you use arrival sequence source files.

To assign unique sequence numbers to all the records, specify the following parameters on the Reorganize Physical File Member (RGZPFM) command:

- KEYFILE(\*NONE), so that the records are not reorganized
- SRCOPT(\*SEQNBR), so that the sequence numbers are changed
- SRCSEQ with a fractional value such as .10 or .01, so that all the sequence numbers are unique

**Note:** Deleted records, if they exist, will be compressed out.

A source file with an arrival sequence access path can be reorganized by sequence number if a logical file for which a keyed sequence access path is specified is created over the physical file.

## Determining When a Source Statement Was Changed

Each source record contains a date field which is automatically updated by SEU if a change occurs to the statement. This can be used to determine when a statement was last changed. Most high-level language compilers print these dates on the compiler lists. The Copy File (CPYF) and Copy Source File (CPYSRCF) commands also print these dates.

Each source member description contains two date and time fields. The first date/time field reflects changes to the member any time it is closed after being updated.

The second date/time field reflects any changes to the member. This includes all changes caused by SEU, commands (such as CRYF and CPYSRCF), authorization changes, and changes to the file status. For example, the FRCRATIO parameter on the Change Physical File (CHGPF) command changes the member status. This date/time field is used by the Save Changed Objects (SAVCHGOBJ) command to determine if the member should be saved. Both date/time fields can be displayed with the Display File Description (DSPFD) command specifying TYPE(\*MBR).

There are two changed date/times shown with source members:

- Last source update date/time. This value reflects any change to the source data records in the member. When a source update occurs, the Last change date/time value is also updated, although there may be a 1- or 2-second difference in that date/time value.
- Last change date/time. This value reflects any changes to the member. This includes all changes caused by SEU, commands (such as CPYF and CPYSRCF), authorization changes, or changes to file status. For example, the FRCRATIO parameter on the CHGPF command changes the member status, and therefore, is reflected in the Last change date/time value.

## Using Source Files for Documentation

You can use the IBM-supplied source file QXTSRC to help you create and update online documentation.

You can create and update QXTSRC members just like any other application (such as QRPGRSRC or QCLSRC) available with SEU. The QXTSRC file is most

useful for narrative documentation, which can be retrieved online or printed. The text that you put in a source member is easy to update by using the SEU add, change, move, copy, and include operations. The entire member can be printed by specifying Yes for the print current source file option on the exit prompt. You can also write a program to print all or part of a source member.

---

## Appendix A. Database File Sizes

The following database file maximums should be kept in mind when designing files on the AS/400 system:

Description	Maximum Value
Number of bytes in a record	32,766 bytes
Number of fields in a record format	8,000 fields
Number of key fields in a file	120 fields
Size of key for physical and logical files	2000 characters <sup>1</sup>
Size of key for ORDER BY (SQL/400) and KEYFLD (OPNQRYF)	10,000 bytes
Number of records contained in a file member	2,147,483,646 records <sup>2</sup>
Number of bytes in a file member	266,757,734,400 bytes <sup>3</sup>
Number of bytes in an access path	4,294,966,272 bytes <sup>3</sup>
Number of keyed logical files built over a physical file member	3,686 files
Number of physical file members in a logical file member	32 members
Number of members that can be joined	32 members
Size of a character or DBCS field	32,766 bytes <sup>4</sup>
Size of a zoned decimal or packed decimal field	31 digits

1 When a first-changed-first-out (FCFO) access path is specified for the file, the maximum value for the size of the key for physical and logical files is 1995 characters.

2 For files with keyed sequence access paths, the maximum number of records in a member varies and can be estimated using the following formula:

$$\frac{2,867,200,000}{10 + (.8 \times \text{key length})}$$

This is an estimated value, the actual maximum number of records can vary significantly from the number determined by this formula.

3 Both the number of bytes in a file member and the number of bytes in an access path must be looked at when message CPF5272 is sent indicating that the maximum system object size has been reached.

4 The maximum size of a variable-length character or DBCS field is 32,740 bytes. DBCS-graphic field lengths are expressed in terms of characters; therefore, the maximums are 16,383 characters (fixed length) and 16,370 characters (variable length).

---

These are maximum values. There are situations where the actual limit you experience will be less than the stated maximum. For example, certain high-level languages can have more restrictive limits than those described above.

Keep in mind that performance can suffer as you approach some of these maximums. For example, the more logical files you have built over a physical file, the greater the chance that system performance can suffer (if you are frequently changing data in the physical file that causes a change in many logical file access paths).

Normally, an AS/400 database file can grow until it reaches the maximum size allowed on the system. The system normally will not allocate all the file space at once. Rather, the system will occasionally allocate additional space as the file grows larger. This method of automatic storage allocation provides the best combination of good performance and effective auxiliary storage space management.

If you want to control the size of the file, the storage allocation, and whether the file should be connected to auxiliary storage, you can use the SIZE, ALLOCATE, and

CONTIG parameters on the Create Physical File (CRTPF) and the Create Source Physical File (CRTSRCPF) commands.

You can use the following formulas to estimate the disk size of your physical and logical files.

- For a physical file (excluding the access path):

$$\text{Disk size} = (\text{number of valid and deleted records} + 1) \times (\text{record length} + 1) + 5120 \times (\text{number of members}) + 1024$$

The size of the physical file depends on the SIZE and ALLOCATE parameters on the CRTPF and CRTSRCPF commands. If you specify ALLOCATE(\*YES), the initial allocation and increment size on the SIZE keyword must be used instead of the number of records.

- For a logical file (excluding the access path):

$$\text{Disk size} = (2560) \times (\text{number of members}) + 1024$$

- For a keyed sequence access path, per member:

$$(\text{number of keys}) \times (\text{key length} + 8) \times (0.8) \times (1.85) + 4096$$

**Notes:**

1. The number of keys is the number of index entries in the access path.
2. The key length is the sum of the length of all the key fields.
3. 0.8 is an estimate of key compression.
4. 1.85 is an estimate of index overhead and space reserved for inserting keys.

This estimate can differ significantly from your file. The keyed sequence access path depends heavily on the data in your records. The only way to get an accurate size is to load your data and display the file description.

The following is a list of minimum file sizes:

Description	Minimum Size
Physical file without a member	1024 bytes
Physical file with a single member	8074 bytes
Keyed sequence access path	9216 bytes

**Note:** Additional space is not required for an arrival sequence access path.

In addition to the file sizes, the system maintains internal formats and directories for database files. (These internal objects are owned by user profile QDBSHR.) The following are estimates of the sizes of those objects:

- For any file not sharing another file's format:

$$\text{Format size} = (96 \times \text{number of fields}) + 512$$

- For files sharing their format with any other file:

$$\text{Format sharing directory size} = (16 \times \text{number of files sharing the format}) + 272$$

- For each physical file and each physical file member having a logical file or logical file member built over it:

$$\text{Data sharing directory size} = (16 \times \text{number of files or members sharing data}) + 272$$

- For each file member having a logical file member sharing its access path:  
Access path sharing directory size = (16 x number of files  
or members sharing access path) + 272



---

## Appendix B. Double-Byte Character Set (DBCS) Considerations

A double-byte character set (DBCS) is a character set that represents each character with 2 bytes. The DBCS supports national languages that contain a large number of unique characters or symbols (the maximum number of characters that can be represented with 1 byte is 256 characters). Examples of such languages include Japanese, Korean, and Chinese.

This appendix describes DBCS considerations as they apply to the database on the AS/400 system.

---

### DBCS Field Data Types

There are two general kinds of DBCS data: bracketed-DBCS data and graphic (nonbracketed) DBCS data. **Bracketed-DBCS data** is preceded by a DBCS shift-out character and followed by a DBCS shift-in character. **Graphic-DBCS data** is not surrounded by shift-out and shift-in characters. The application program might require special processing to handle bracketed-DBCS data that would not be required for graphic-DBCS data.

The specific DBCS data types (specified in position 35 on the *DDS Coding Form*) are:

Entry	Meaning
O	DBCS-open: A character string that contains both single-byte and bracketed double-byte data.
E	DBCS-either: A character string that contains either all single-byte data or all bracketed double-byte data.
J	DBCS-only: A character string that contains only bracketed double-byte data.
G	DBCS-graphic: A character string that contains only nonbracketed double-byte data.

**Note:** Files containing DBCS data types can be created on a single-byte character set (SBCS) system. Files containing DBCS data types can be opened and used on a SBCS system, however, coded character set identifier (CCSID) conversion errors can occur when the system tries to convert from a DBCS or mixed CCSID to a SBCS CCSID. These errors will not occur if the job CCSID is 65535.

### DBCS Constants

A constant identifies the actual character string to be used. The character string is enclosed in apostrophes and a string of DBCS characters is surrounded by the DBCS shift-out and shift-in characters (represented by the characters < and > in the following examples). A DBCS-graphic constant is preceded by the character G. The types of DBCS constants are:

Type	Example
DBCS-Only	'<A1A2A3>'

DBCS-Open	'<A1A2A3>BCD'
DBCS-Graphic	G'<A1A2A3>'

## DBCS Field Mapping Considerations

The following chart shows what types of data mapping are valid between physical and logical files for DBCS fields:

Physical File Data Type	Logical File Data Type					
	Character	Hexadecimal	DBCS-Open	DBCS-Either	DBCS-Only	DBCS-Graphic
Character	Valid	Valid	Valid	Valid	Not valid	Not valid
Hexadecimal	Valid	Valid	Valid	Valid	Valid	Not valid
DBCS-open	Not valid	Valid	Valid	Not valid	Not valid	Not valid
DBCS-either	Not valid	Valid	Valid	Valid	Not valid	Not valid
DBCS-only	Not valid	Valid	Valid	Valid	Valid	Valid
DBCS-graphic	Not valid	Not valid	Valid	Valid	Valid	Valid

## DBCS Field Concatenation

When fields are concatenated, the data types can change (the resulting data type is automatically determined by the system).

- OS/400 assigns the data type based on the data types of the fields that are being concatenated. When DBCS fields are included in a concatenation, the general rules are:
  - If the concatenation contains one or more hexadecimal (H) fields, the resulting data type is hexadecimal (H).
  - If all fields in the concatenation are DBCS-only (J), the resulting data type is DBCS-only (J).
  - If the concatenation contains one or more DBCS (O, E, J) fields, but no hexadecimal (H) fields, the resulting data type is DBCS open (O).
  - If the concatenation contains two or more DBCS open (O) fields, the resulting data type is a variable-length DBCS open (O) field.
  - If the concatenation contains one or more variable-length fields of any data type, the resulting data type is variable length.
  - A DBCS-graphic (G) field can be concatenated only to another DBCS-graphic field. The resulting data type is DBCS-graphic (G).
- The maximum length of a concatenated field varies depending on the data type of the concatenated field and length of the fields being concatenated. If the concatenated field is zoned decimal (S), its total length cannot exceed 31 bytes. If the concatenated field is character (A), DBCS-open (O), or DBCS-only (J), its total length cannot exceed 32,766 bytes (32,740 bytes if the field is variable length).

The length of DBCS-graphic (G) fields is expressed as the number of double-byte *characters* (the actual length is twice the number of characters); therefore, the total length of the concatenated field cannot exceed 16,383 characters (16,370 characters if the field is variable length).



- In join logical files, the fields to be concatenated must be from the same physical file. The first field specified on the CONCAT keyword identifies which physical file is used. The first field must, therefore, be unique among the physical files on which the logical file is based, or you must also specify the JREF keyword to specify which physical file to use.
- The use of a concatenated field must be I (input only).
- REFSHIFT cannot be specified on a concatenated field that has been assigned a data type of O or J.

**Notes:**

1. When bracketed-DBCS fields are concatenated, a shift-in at the end of one field and a shift-out at the beginning of the next field are removed. If the concatenation contains one or more hexadecimal fields, the shift-in and shift-out pairs are only eliminated for DBCS fields that precede the first hexadecimal field.
2. A concatenated field that contains DBCS fields must be an input-only field.
3. Resulting data types for concatenated DBCS fields may differ when using The Open Query File (OPNQRYP) command. See “Using Concatenation with DBCS Fields through OPNQRYP” on page B-5 for general rules when DBCS fields are included in a concatenation.

---

## DBCS Field Substring Operations

A substring operation allows you to use part of a field or constant in a logical file. For bracketed-DBCS data types, the starting position and the length of the substring refer to the number of *bytes*; therefore, each double-byte character counts as two positions. For the DBCS-graphic (G) data type, the starting position and the length of the substring refer to the number of *characters*; therefore, each double-byte character counts as one position.

---

## Comparing DBCS Fields in a Logical File

When comparing two fields or a field and constants, fixed-length fields can be compared to variable-length fields as long as the types are compatible. Figure B-1 describes valid comparisons for DBCS fields in a logical file.

Figure B-1 (Page 1 of 2). Valid Comparisons for DBCS Fields in a Logical File

	Any Numeric	Character	Hexadecimal	DBCS-Open	DBCS-Either	DBCS-Only	DBCS-Graphic	Date	Time	Time Stamp
<b>Any Numeric</b>	Valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid
<b>Character</b>	Not valid	Valid	Valid	Valid	Valid	Not valid	Not valid	Not valid	Not valid	Not valid
<b>Hexadecimal</b>	Not valid	Valid	Valid	Valid	Valid	Valid	Not valid	Not valid	Not valid	Not valid
<b>DBCS-Open</b>	Not valid	Valid	Valid	Valid	Valid	Valid	Not valid	Not valid	Not valid	Not valid
<b>DBCS-Either</b>	Not valid	Valid	Valid	Valid	Valid	Valid	Not valid	Not valid	Not valid	Not valid

Figure B-1 (Page 2 of 2). Valid Comparisons for DBCS Fields in a Logical File

	Any Numeric	Character	Hexadecimal	DBCS-Open	DBCS-Either	DBCS-Only	DBCS-Graphic	Date	Time	Time Stamp
<b>DBCS-Only</b>	Not valid	Not valid	Valid	Valid	Valid	Valid	Not valid	Not valid	Not valid	Not valid
<b>DBCS-Graphic</b>	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Valid	Not valid	Not valid	Not valid
<b>Date</b>	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Valid	Not valid	Not valid
<b>Time</b>	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Valid	Not valid
<b>Time Stamp</b>	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Valid

## Using DBCS Fields in the Open Query File (OPNQRYF) Command

This section describes considerations when using DBCS fields in the Open Query File (OPNQRYF) command.

### Using the Wildcard Function with DBCS Fields

Use of the wildcard (%WLDCRD) function with a DBCS field differs depending on whether the function is used with a bracketed-DBCS field or a DBCS-graphic field.

When using the wildcard function with a bracketed-DBCS field, both single-byte and double-byte wildcard values (asterisk and underline) are allowed. The following special rules apply:

- A single-byte underline refers to one EBCDIC character; a double-byte underline refers to one double-byte character.
- A single- or double-byte asterisk refers to any number of characters of any type.

When using the wildcard function with a DBCS-graphic field, only double-byte wildcard values (asterisk and underline) are allowed. The following special rules apply:

- A double-byte underline refers to one double-byte character.
- A double-byte asterisk refers to any number of double-byte characters.

## Comparing DBCS Fields Through OPNQRYF

When comparing two fields or constants, fixed length fields can be compared to variable length fields as long as the types are compatible. Figure B-2 describes valid comparisons for DBCS fields through the OPNQRYF command.

Figure B-2 (Page 1 of 2). Valid Comparisons for DBCS Fields through the OPNQRYF Command

	Any Numeric	Character	Hexadecimal	DBCS-Open	DBCS-Either	DBCS-Only	DBCS-Graphic	Date	Time	Time Stamp
<b>Any Numeric</b>	Valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid
<b>Character</b>	Not valid	Valid	Valid	Valid	Valid	Not valid	Not valid	Valid	Valid	Valid

Figure B-2 (Page 2 of 2). Valid Comparisons for DBCS Fields through the OPNQRYF Command

	Any Numeric	Char- acter	Hexa- decimal	DBCS- Open	DBCS- Either	DBCS- Only	DBCS- Graphic	Date	Time	Time Stamp
<b>Hexa- decimal</b>	Not valid	Valid	Valid	Valid	Valid	Valid	Not valid	Valid	Valid	Valid
<b>DBCS- Open</b>	Not valid	Valid	Valid	Valid	Valid	Valid	Not valid	Valid	Valid	Valid
<b>DBCS- Either</b>	Not valid	Valid	Valid	Valid	Valid	Valid	Not valid	Valid	Valid	Valid
<b>DBCS- Only</b>	Not valid	Not valid	Valid	Valid	Valid	Valid	Not valid	Not valid	Not valid	Not valid
<b>DBCS- Graphic</b>	Not valid	Not valid	Not valid	Not valid	Not valid	Not valid	Valid	Not valid	Not valid	Not valid
<b>Date</b>	Not valid	Valid	Valid	Valid	Valid	Not valid	Not valid	Valid	Not valid	Not valid
<b>Time</b>	Not valid	Valid	Valid	Valid	Valid	Not valid	Not valid	Not valid	Valid	Not valid
<b>Time Stamp</b>	Not valid	Valid	Valid	Valid	Valid	Not valid	Not valid	Not valid	Not valid	Valid

## Using Concatenation with DBCS Fields through OPNQRYF

When using the Open Query File (OPNQRYF) concatenation function, the OS/400 program assigns the resulting data type based on the data types of the fields being concatenated. When DBCS fields are included in a concatenation, the resulting data type is generally the same as concatenated fields in a logical file, with some slight variations. The following rules apply:

- If the concatenation contains one or more hexadecimal (H) fields, the resulting data type is hexadecimal (H).
- If all fields in the concatenation are DBCS-only (J), the resulting data type is variable length DBCS-only (J).
- If the concatenation contains one or more DBCS (O, E, J) fields, but no hexadecimal (H) fields, the resulting data type is variable length DBCS open (O).
- If the concatenation contains one or more variable length fields of any data type, the resulting data type is variable length.
- A DBCS-graphic (G) field can be concatenated only to another DBCS-graphic field. The resulting data type is DBCS-graphic (G).

## Using Sort Sequence with DBCS

When a sort sequence is specified, no translation of the DBCS data is done. Only SBCS data in DBCS-either or DBCS-open fields is translated.



## Appendix C. Database Lock Considerations

Figure C-1 summarizes some of the most commonly used database functions and the types of locks they place on database files. The types of locks are explained on the next page.

Figure C-1 (Page 1 of 2). Database Functions and Locks

Function	Command	File Lock	Member/Data Lock	Access Path Lock
Add Member	ADDPFM, ADDLFM	*EXCLRD		*EXCLRD
Change File Attributes	CHGPF, CHGLF	*EXCL	*EXCLRD	*EXCLRD
Change Member Attributes	CHGPFM, CHGLFM	*SHRRD	*EXCLRD	
Change Object Owner	CHGOBJOWN	*EXCL		
Check Object	CHKOBJ	*SHRNUPD		
Clear Physical File Member	CLRPFM	*SHRRD	*EXCLRD <sup>3</sup>	
Create Duplicate Object	CRTDUPOBJ	*EXCL (new object) *SHRNUPD (object)		
Create File	CRTPF, CRTLf, CRTSRCPF	*EXCL		
Delete File	DLTF	*EXCL		*EXCLRD
Grant/Revoke Authority	GRTOBJAUT, RVKOBJAUT	*EXCL		
Initialize Physical File Member	INZPFM	*SHRRD	*EXCLRD	
Move Object	MOV OBJ	*EXCL		
Open File	OPNDBF, OPNQRYF	*SHRRD	*SHRRD	*EXCLRD
Rebuild Access Path	EDTRBDAP, OPNDBF	*SHRRD	*SHRRD	*EXCLRD
Remove Member	RMVM	*EXCLRD	*EXCL	*EXCLRD
Rename File	RNMOBJ	*EXCL	*EXCL	*EXCL
Rename Member	RNMM	*EXCLRD	*EXCL	*EXCL
Reorganize Physical File Member	RGZPFM	*SHRRD	*EXCL	
Restore File	RSTLIB, RSTOBJ	*EXCL		
Save File	SAVLIB, SAVOBJ, SAVCHGOBJ	*SHRNUPD <sup>1</sup>	*SHRNUPD <sup>2</sup>	

Figure C-1 (Page 2 of 2). Database Functions and Locks

Function	Command	File Lock	Member/Data Lock	Access Path Lock
1	For save-while-active, the file lock is *SHRUPD initially, and then the lock is reduced to *SHRRD. See the <i>Advanced Backup and Recovery Guide</i> for a description of save-while-active locks for the save commands.			
2	For save-while-active, the member/data lock is *SHRRD.			
3	The clear does not happen if the member is open in this process or any other process.			

The following table shows the valid lock combinations:

Lock	*EXCL	*EXCLRD	*SHRUPD	*SHRNUPD	*SHRRD
*EXCL <sup>1</sup>					
*EXCLRD <sup>2</sup>					X
*SHRUPD <sup>3</sup>			X		X
*SHRNUPD <sup>4</sup>				X	X
*SHRRD <sup>5</sup>		X	X	X	X

- 1 Exclusive lock (\*EXCL). The object is allocated for the exclusive use of the requesting job; no other job can use the object.
- 2 Exclusive lock, allow read (\*EXCLRD). The object is allocated to the job that requested it, but other jobs can read the object.
- 3 Shared lock, allow read and update (\*SHRUPD). The object can be shared either for read or change with other jobs.
- 4 Shared lock, read only (\*SHRNUPD). The object can be shared for read with other jobs.
- 5 Shared lock (\*SHRRD). The object can be shared with another job if the job does not request exclusive use of the object.

---

# Appendix D. Design Guidelines for OPNQRYP Performance

---

## Overview

This chapter discusses many sections related to programming for optimizing performance of a query application. These sections not only apply to OPNQRYP, but are also a good guide for other products that use the AS/400 database techniques, such as the SQL/400 program, the Query/400 program, and Query Management/400.

The following sections are discussed:

1. Data management methods

This covers the algorithms used to retrieve data from the disk and includes a discussion of access paths and row selection techniques.

2. Optimizer

The optimizer is the function that decides how to gather data which should be returned to the program. This section covers the techniques and rules employed by the optimizer for performing this task including cost estimating, access plan validation, and join optimization.

3. Optimizer messages

4. Tips and techniques

Before you start designing for performance you should think about the following:

1. Consider performance:

- If there are over 10,000 rows, the effect on performance is *noticeable*.
- If there are over 100,000 rows, the effect on performance is a *concern*.
- If Complex queries are used repetitively
- If there are multiple work stations with high transaction rates

2. Optimize resources:

- I/O Usage
- CPU usage
- Effective usage of indexes
- Concurrency (COMMIT)

**Note:** The information in this chapter is complex. It may be helpful to experiment with and to verify some of the information using an AS/400 system as you read this chapter.

## Definition of Terms

Understanding the following terms is necessary in order to understand the information in this chapter:

### Access plan

A control structure that describes the actions necessary to satisfy each query request. An access plan contains information about the data and how to extract it.

**Binding** The establishment of a relationship between the program or query and any specified file.

**Cursor** The cursor is a pointer. It references the record to be processed.

**Open data path (ODP)**

An I/O area that is opened for every file during the processing of a high-level language (HLL) program.

**Process access group (PAG)**

The ODP is kept in a structure that is unique to each job that is processing and is called a process access group (PAG). The size of this PAG can vary greatly depending on the number of ODPs and the number of active programs. Each program has work areas (for many HLLs these are known as program static storage areas or PSSAs) that are also contained in the PAG. While a job is processing, the system has to manage this PAG. This results in system overhead. The larger the PAG, the greater this overhead can be. It is reflected through a larger program working set.

## OS/400 Query Component

Traditional HLL program I/O statements access data one record at a time. You can use I/O statements in conjunction with logical files to provide relational operations such as:

- Record selection
- Sequence
- Join
- Project

This is often the most efficient manner for data retrieval.

Query products are best used when logical files are not available for data retrieval requests, or if functions are required that logical files cannot support, are too difficult to write, or would perform poorly, for example, the Distinct, Group by, Subquery, and Like functions.

The query products use something more sophisticated to perform these functions. It is done with the access plan in combination with a specialized, high-function query routine called the OS/400 query component, which is internal to the OS/400 program. (The OS/400 query component should not be confused with the Query/400 licensed program.) The advantage of this function is that, because the query requests are created at run time, there are often fewer permanent access paths than are required for multiple logical files.

Nine programs and functions on the AS/400 system use this OS/400 query component:

- OPNQRYF
- SQL/400 run-time support
- Query/400 run-time support
- Query/38 run-time support
- PC Support file transfer
- Query Management/400
- OfficeVision/400\* (for document searches)
- Performance Tools (for report generation)
- Q&A Database



The difference between the terminology of SQL/400, Query/400, and Query/38 versus SQL/400 *run-time support*, Query/400 *run-time support*, and Query/38 *run-time support* is that the former group refers to the names of the licensed program. The licensed programs are not required to run each of these, as the run-time support comes with the OS/400 program and not the licensed programs.

Figure D-1 helps to explain the relationship between the various functions that use the OS/400 query component at run time, and the way in which traditional HLL program I/O requests are satisfied. Note that HLL program I/O requests go directly to the database support to retrieve the data. Query product requests call on the OS/400 query component, which uses the optimizer before calling the database support to create the ODP to the data. Once an ODP has been created, no difference exists between HLL I/O requests and the I/O requests of these query products. Both send requests for data to the database support.

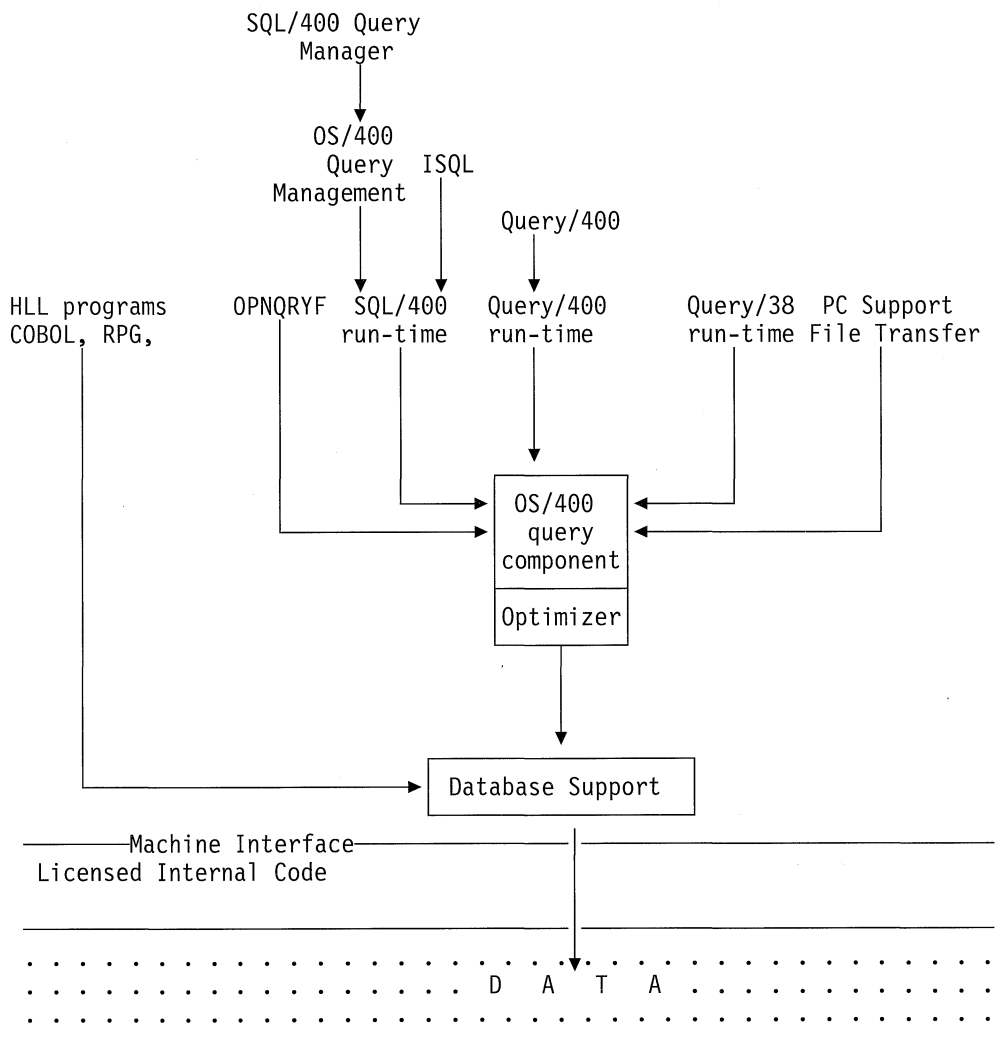


Figure D-1. Methods of Accessing AS/400 Data

**The Optimizer:** The optimizer is a major part of the OS/400 query component. It makes decisions to improve the performance of a query. In some queries, the optimizer may decide to build a temporary index over the data. When this happens, the time required for the index build delays the query processing. The database contains information such as which indexes are already built on a file. If necessary,

the optimizer uses these indexes to assist performance. The main objective of query performance analysis is the correct use of indexes. The programmer is responsible for creating indexes that the optimizer will use.

A more detailed discussion on the optimizer, access plans, and the QDT appears later in this chapter. See “The Optimizer” on page D-10 for more information.

---

## Data Management Methods

AS/400 data management provides various methods to retrieve data. This section introduces the fundamental techniques used in the OS/400 program and the Licensed Internal Code. These methods or combinations of methods are used by the OS/400 query component to access the data.

For complex query tasks, you can find different solutions that satisfy your requirements for retrieval of data from the database. This appendix is not a cookbook that helps to find the best performing variation for a query. You have to understand enough about the creation of the access plan and the decisions of the optimizer (discussed in “The Optimizer” on page D-10) to find the solution that suits your needs. For this reason, this section discusses the following topics that are fundamental for data retrieval from the AS/400 database:

- Access path
- Access method

### Access Path

An access path is:

- The order in which records in a database file are organized for processing.
- The path used to locate data specified in a query. An access path can be indexed, sequential, or a combination of both.

#### Arrival Sequence Access Path

An arrival sequence access path is the order of records as they are stored in the file. Processing files using the arrival sequence access path is similar to processing sequential or direct files on traditional systems.

#### Keyed Sequence Access Path

A keyed sequence access path provides access to a database file that is arranged according to the contents of key fields (indexes). The keyed sequence is the order in which records are retrieved. The access path is automatically maintained whenever records are added to or deleted from the file, or whenever the contents of the index fields are changed. The best example of a keyed sequence access path is a logical file (created using the CRTLF command).

Fields that are good candidates for creating keyed sequence access paths are:

- Those frequently referenced in query record selection predicates (QRYSLT parameter in OPNQRYP command).
- Those frequently referenced in grouping or ordering specifications (GRPFLD or KEYFLD parameters in OPNQRYP command).
- Those used to join files (see “Join Optimization” on page D-13).

For a further description of access paths, refer to the *Data Management Guide*, SC41-9658.

## Access Method

The use of access methods is divided between the Licensed Internal Code and the OS/400 query component. The Licensed Internal Code does the low-level processing: selection, join functions, and access path creation. These low-level functions actually involve reading and checking the data. Records that meet the selection criteria are passed back to the calling program. (See Figure D-1 on page D-3 for an illustration.)

The query optimization process chooses the most efficient access method for each query and keeps this information in the access plan. The type of access is dependent on the number of records, the number of page faults<sup>1</sup>, and other criteria (refer to “The Optimizer” on page D-10).

This section discusses the possible methods the optimizer can use to retrieve data. The general approach is to either do a data scan (defined below), use an existing index, create a temporary index from the data space, create a temporary index from an existing index, or use the query sort routine. Selection can be implemented through:

- Dynamic processing method
- Key selection method
- Key positioning method
- Index-from-index method

Definition of terms used in the following section:

- The internal object that contains the data in a file is referred to as a **data space**.
- The first key field of an index over multiple fields is referred to as the **primary** or **left-most** key.

### Dynamic Processing Access Method

The records in the file are processed in no guaranteed order. They will be in arrival sequence. If you want the result in a particular sequence, you must specify the KEYFLD parameter in the OPNQRYF command. Because indexes are not used in arrival sequence order, all records in the file are read. This operation is referred to as a *data space scan*. The selection criteria is applied to each record, and only the records that match the criteria are returned to the calling application.

Dynamic processing can be very efficient for the following reasons:

- It minimizes the number of page I/O operations because all records in a given page are processed, and once the page has been retrieved, it will not be retrieved again.
- The database manager can easily predict the sequence of pages from the data space for retrieval; therefore, it can schedule asynchronous I/O of the pages into main storage from auxiliary storage (commonly referred to as *pre-fetching*).

---

<sup>1</sup> An interrupt that occurs when a program refers to a (512 byte) page that is not in main storage.

The idea is that the page would be available in main storage by the time the database manager needs to examine the data.

This selection method is very good when a large percentage of the records are to be selected (greater than approximately 20%).

Dynamic processing can be adversely affected when selecting records from a file containing deleted records. As you may recall, the delete operation only marks records as deleted. For dynamic processing, the database manager is going to read all of the deleted records, even though none will be selected. You should use the Reorganize Physical File Member (RGZPFM) CL command to eliminate deleted records.

Dynamic processing is not very efficient when a small number of records will be selected. Because all records in the file are examined, this leads to consumption of wasted I/O and processing unit resources.

### **Key Selection Access Method**

This access method requires keyed sequence access paths. The entire index is read and any selection criteria that references the key fields of the index is applied against the index. The advantage of this method is that the data space is only accessed to retrieve records that satisfy the selection criteria applied against the index. Any selection not performed through this key selection method is performed using dynamic processing at the data-space level.

The key selection access method can be very expensive if the search condition applies to a large number of records because:

- The whole index is processed.
- For every key read from the index, a random I/O to the data space occurs.

Normally, the optimizer would choose to use dynamic processing when the search condition applies to a large number of records. Only if ordering, grouping, or join operations were specified (these options force the use of an index) would the optimizer choose to use key selection when the search condition selects more than approximately 20% of the keys. In these cases, the optimizer may choose to create a temporary index rather than use an existing index. When the optimizer creates a temporary index it uses a 16K page size. An index created using the CRTLF command uses only a 2K page size. The optimizer also processes as much selection as possible when building the temporary index. Therefore, nearly all temporary indexes built by the optimizer are select/omit or sparse indexes. The page size difference and corresponding performance improvement from swapping in fewer pages may be enough to overcome the overhead of creating an index.

Before Version 2 Release 2, if ordering was specified, the optimizer was forced to use an index and, consequently, the key selection access method. Since Version 2 Release 2, the same query could be resolved using the query sort routine if you specified the following parameters on the OPNQRYF command:

- ALWCPYDTA(\*OPTIMIZE) and COMMIT(\*NO)
- ALWCPYDTA(\*OPTIMIZE) and COMMIT(\*YES) and the journal is opened with a commit level of \*NONE, \*CHG, or \*CS

## Key Positioning Access Method

This access method is very similar to the key selection access method. They both require a keyed sequence access path. Unlike key selection access method, where processing starts at the beginning of the index and continues to the end, key positioning access method uses selection against the index to position directly to a range of keys that match some, or all, of the selection criteria. It reads all the keys from this range and performs any remaining key selection, similar to the selection performed by the key selection method. Any selection not performed through key positioning or key selection is performed as dynamic processing at the data-space level. Because key positioning only processes a subset of the keys in the index, it is a better performing method than key selection.

The key positioning method is most efficient when a small percentage of records are to be selected (less than approximately 20%). If more than approximately 20% of the records are to be selected, the optimizer generally chooses to:

- Use dynamic processing (index is not required)
- Use key selection (if an index is required)
- Use query sort routine (if conditions apply)

For queries that do not require an index (no ordering, grouping, or join operations), the optimizer tries to find an existing index to perform key positioning. If no existing index can be found, the optimizer stops trying to use keyed access to the data because it is faster to use dynamic processing than it is to build an index and then perform key positioning.

The following example illustrates a query where the optimizer could choose the key positioning method:

Create an access path INDEX1 keyed over field FIELD1
OPNQRYF FILE((INDEX1)) QRYSLT('FIELD1 *EQ 'C''')

In this example, the database support uses *Index1* to position to the first index entry with *FIELD1* value equal to 'C'. For each key equal to 'C', it randomly accesses the data space (random accessing occurs because the keys may not be in the same sequence as the records in the data space) and selects the record. If additional selection needs to be applied to the record (if additional selections exist that did not match the primary key of the index), they are performed while selecting the records from the data space. The query ends when the key selection moves beyond the key value of 'C'.

Note that for this example all index entries processed and records retrieved meet the selection criteria. If additional selections are added that cannot be performed through key positioning (field does not match the primary key of the index), the optimizer uses key selection to perform as much selection as possible. The remaining selection is performed as dynamic processing at the data-space level.

The key positioning access method has additional processing capabilities. One such capability is to perform range selection across more than one value. For example:

```
OPNQRYF FILE((ANYFILE))
        QRYSLT('FIELD1 *EQ %RANGE(''C'' ''D'')')
```

In this example, the selection is positioned to the first index entry equal to value 'C' and then processes records until the last index entry for 'D' is processed.

A further extension of this access method, called *multirange* key positioning, is available. It allows for the selection of rows for multiple ranges of values for a given primary key:

```
OPNQRYF FILE((ANYFILE))
        QRYSLT('FIELD1 *EQ %RANGE(''C'' ''D'')
              *OR 'FIELD1 *EQ %RANGE(''F'' ''H'')')
```

In this example, the positioning and processing technique is used twice, once for each range of values.

**Note:** Key positioning works only with the primary key field. The efficiency of the key positioning access method depends on how selective the primary key of the index is.

### Index-From-Index Access Method

The database manager can build a temporary index from an existing index without having to read all of the records in the data space. Generally speaking, this selection method is one of the most efficient. The temporary index that is created only contains keys for records that meet the selection predicates, similar to a select/omit or sparse index. The optimizer chooses this step when:

- The query requires an index because it uses grouping, ordering, or join processing.
- A permanent index exists that has a selection field as the primary key and the primary key is very selective.
- The selection field is not the same as the order by, group by, or join to field.

To process using the index-from-index access method, the database manager first uses key positioning on the permanent index with the query selection criteria. Secondly, selected record entries are used to build index entries in the new temporary index. The result is an index containing entries in the required key sequence for records that match the selection criteria. Since Version 2 Release 2, your query may be resolved using the query sort routine if you specify either:

- ALWCOPYDTA(\*OPTIMIZE) and COMMIT(\*NO)
- ALWCOPYDTA(\*OPTIMIZE) and COMMIT(\*YES) and the journal is opened with a commit level of \*NONE, or \*CHG, or \*CS

This decision is based on the number of records to be retrieved.

A common index-from-index access method example:

```

Create an access path INDEX1
keyed over field FIELD1

OPNQRYF FILE((INDEX1))
        QRYSLT('FIELD1 *EQ 'C''')
        KEYFLD((FIELD2))

```

For this example, a temporary select/omit access path with primary key FIELD2 is created, containing index entries for those records where *FIELD1* = 'C', assuming *FIELD1* = 'C' selects fewer than approximately 20% of the records. The following table represents a summary of the data management methods discussed.

Figure D-2. Summary of Data Management Methods

Access Method	Selection Process	Good When	Not Good When	Selected When	Advantages
Dynamic processing	Reads all records. Selection criteria applied to data in data space.	> 20% records selected	< 20% records selected	No ordering, grouping, or joining and > 20% records selected.	Minimizes page I/O through pre-fetching.
Key selection	Selection criteria applied to index.	Ordering, grouping and joining	Large number of records selected.	Index is required and cannot use key positioning method.	Data space accessed only for records matching key selection criteria.
Key positioning	Selection criteria applied to range of index entries. Commonly used option.	< 20% records are selected	> 20% records are selected	Selection field matches primary key and selects < 20% records.	Index and Data space accessed only for records matching selection criteria.
Index-from-index	Key row positioning on permanent index. Builds temporary index over selected index entries.	Ordering, grouping and join operations	> 20% records are selected	No existing index to satisfy ordering but existing index does satisfy selection and selecting < 20% records.	Index and data space accessed only for records matching selection criteria.
Sort routine	Order data read using dynamic processing or key positioning.	> 20% records selected or large result set of records	< 20% records selected or small result set of records	Ordering specified; either no index exists to satisfy the ordering or a large result set is expected	See dynamic processing and key positioning in this table.

---

## The Optimizer

The optimizer is an important module of the OS/400 query component because it makes the key decisions for good database performance. Its main objective is to find the most efficient access method to the data. This section discusses how the optimizer works in general. The exact algorithms are too complex to be described in detail here and are subject to change from release to release.

Query optimization is a trade-off between the time spent to select a query implementation and the time spent to process it. Query optimization must have:

- A quick interactive response
- An efficient use of total machine resources

In deciding how to access data, the optimizer:

- Determines possible uses
- Picks the optimal use for the OS/400 query component to process

## Implementation Cost Estimation

At run time, the optimizer chooses an optimal use of the query by calculating an *implementation cost* given the current state of the database. The optimizer models the access cost of each of the following:

- Reading records directly from the file (dynamic processing)
- Reading records through an access path (can use either key selection or key positioning)
- Creating an access path directly from the data space (builds a temporary access path, and then can use any selection option on the temporary access path)
- Creating an access path from an existing access path (index-from-index)
- Using the query sort routine (if conditions are satisfied)

The cost of a particular process is the sum of:

- The start-up cost
- The cost associated with the given optimize parameter (\*FIRSTIO, \*ALLIO or \*MINWAIT)

**\*FIRSTIO** Minimize the time required to retrieve the first buffer of records from the file. Biases the optimization toward not creating an index. Either a data scan or an existing index is preferred.

When \*FIRSTIO is selected, users may also pass in the number of records they expect to retrieve from the query. The optimizer uses this value to determine the percentage of records that will be returned and optimizes accordingly. A small value would minimize the time required to retrieve the first n records, similar to \*FIRSTIO. A large value would minimize the time to retrieve all n records, similar to \*ALLIO.

**\*ALLIO** Minimize the time to process the whole query assuming that all query records are read from the file. Does not bias the optimizer to any particular access method.



**Note:** If you specify ALWCPYDTA(\*OPTIMIZE) and use the sort routine, your query resolves according to the \*ALLIO optimize parameter.

**\*MINWAIT** Minimize delays when reading records from the file. Minimize I/O time at the expense of open time. Biases optimization toward either creating a temporary index or performing a sort. Either an index is created or an existing index used.

- The cost of any access path creations
- The expected number of page faults to read the records
- The expected number of records to process

Page faults and number of records processed may be predicted by:

- Statistics the optimizer can obtain from the database objects, including:
  - File size
  - Record size
  - Index size
  - Key size
- A weighted measure of the expected number of records to process based on what the relational operators in the record selection predicates (called *default filter factors*) are likely to retrieve:
  - 10% for equal
  - 33% for less-than, greater-than, less-than-equal-to, or greater-than-equal-to
  - 90% for not equal
  - 25% for RANGE
  - 10% for each VALUES value

Key range estimates are based on the primary (left-most) key of the existing indexes. Generally, six to seven indexes are estimated during the optimization process. The default filter factors may then be further refined by the estimate based on the key range. If an index exists whose primary key matches a field used in a record selection predicate, it can be used to estimate the number of keys in that index matching the selection criteria. The estimate of the number of keys is based on the number of pages and key density of the machine index without actually accessing the keys. Full indexes over fields used in selection predicates can significantly help optimization.

Page faults and the number of records processed are dependent on the type of access the optimizer chooses. Refer to “Data Management Methods” on page D-4 for more information on access methods. For queries that do not require an index, access can be done using dynamic processing or reading through an existing index (key positioning). For queries that require an index (ordering, grouping, or joining is specified), the optimizer may decide to:

- Use an existing index
- Create a new index
- Create an index from another index
- Use a sort routine (for ordering only)

## Access Plan and Validation

For any query, whenever optimization occurs, an optimized plan of how to access the requested data is developed. The information is kept in what is called a miniplan. The miniplan, along with the query definition template (QDT) that is used to interface with the optimizer, make up an access plan. For OPNQRYF, an access plan is created but is not saved. A new access plan is created each time the OPNQRYF command is processed. For Query/400, an access plan is saved as part of the query definition object. For the SQL/400 program, the access plan is saved as part (in the *associated space*) of the program containing embedded SQL statements.

## Optimizer Decision-Making Rules

In performing its function, the optimizer works by a general set of guidelines in choosing the best method for accessing data. The general strategies of this process are:

- Determine the (default) filter factor for each predicate in the selection clause
- Extract attributes of the file from internally stored information
- If an index is required, determine the cost of creating an index over the file or the cost of using a sort routine (if conditions apply)
- If an index is not required, determine the cost of dynamic processing
- For each index available, in the order of most recently created to oldest, and while the optimizer has not spent too much time optimizing:
  - Extract attributes of the index from internally stored statistics.
  - Determine if the index meets the selection criteria.
  - If a selection predicate matches the primary key of the index, perform estimate key range to determine the true filter factor of the predicate.
  - Determine the cost of using the index. Use the estimated page faults and the predicate filter factors to help determine the cost.
  - Compare the cost of using this index with the previous cost (current best).
  - Pick the cheaper one.
  - Continue search for best index until time out or no more indexes.

The *have not spent too much time* factor controls how much time is spent picking an implementation. It is based on how much time has been spent so far versus the current best implementation found.

For small files, the OS/400 query component spends little time in query optimization. It arbitrarily picks one of the first good implementations it finds. For large files, the OS/400 query component considers more of the indexes. Generally, the optimizer considers five or six indexes (for each file of a join) before running out of optimization time.

Since Version 2 Release 1.1, if you specify OPTALLAP(\*YES), the optimizer does not time out and considers all indexes during its optimization phase.

## Join Optimization

Join is a complex function that requires special attention in order to get good performance. For all join operations, the OS/400 database manager requires the use of an access path over the secondary files in the join. If no usable access path exists, the OS/400 database manager builds it. A keyed sequence access path (index) is not required for the primary file unless sort fields are selected from this file, and only from this file. It is important to create indexes to match frequently used join selection criteria. The index should match the join-to fields selected from the secondary file.

```
OPNQRYF FILE((FILE1) (FILE2))
        FORMAT(FILE1)
        JFLD((FILE1/FLDA FILE2/FLDX))
```

For the above example, an access path would be required over FILE2/FLDX. Note that for an inner join (JDFTVAL(\*NO)), the optimizer may decide to switch the join order of the files. If this occurs, an access path would be required over FILE1/FLDA.

Be as selective as possible on all files joining to narrow the number of records that will result from the join operation. This can significantly reduce the amount of I/O required to run the query.

Avoid joining files without a JFLD or QRYSLT clause; otherwise, the result would be that *all* of the records in one file would be joined to *every* record in the other files. This type of operation could result in a large amount of I/O and may affect the overall system performance. The result file of nonequal joins could also contain a large number of records.

OS/400 database manager attempts to order the files from *smallest* to *largest*, depending on the estimated number of records returned. Generally, the query runs more efficiently when the files are ordered in this way.

The following join ordering algorithm is used to determine the order of the files:

1. Determine an access mode for each individual file. Estimate the number of records returned for each file.
2. Determine a cost for each join combination based mainly on the expected number of records returned after the join.

The join order combinations estimated for a four file join would be:

1-2 2-1 1-3 3-1 1-4 4-1 2-3 3-2 2-4 4-2 3-4 4-3

3. Choose the combination with the lowest cost. The lowest cost determines the primary and first secondary file (for example, 2 3 x x).
4. Determine the cost, for each remaining file, of joining to the first secondary file (for example, 3-1 3-4).
5. Choose the next secondary file. This is the file with the lowest join cost when joined to the first secondary file (for example, 2 3 1 x).
6. Repeat the join cost calculation until the full join order has been determined (for example, 2 3 1 4).

**Note:** Specifying KEYFLD for only one of the files always makes this file the primary file for the join operation.

Specifying KEYFLD for only one of the files can be useful when trying to determine which file should be used as the primary file. Note that if you use the KEYFLD on a field from the *larger* file, the larger file becomes the primary file and this may result in less efficient processing than if the *smallest* file were the primary file.

When the sort routine is used (ALWCPYDTA(\*YES) specified), the optimizer is free to switch the join order of the files. Specifying a KEYFLD for only one of the files does not make that file the primary file when the sort routine is used.

See “Join Optimization” on page D-20 for more detail on this and for other join performance tips.

---

## Optimizer Messages

A feature available with OS/400 Version 2 Release 1 is the job log messages sent by the query optimizer. For Version 1 Release 3, the messages are available through PTF. You may order the following PTFs:

- PTF SF05773 and SF05772 for the SQL/400 program and OPNQRYF command
- PTF SF05780 and SF05779 for the SQL/400 program only

The query optimizer provides you with information messages on the current query processing when the job is under debug mode. These messages appear for OPNQRYF, SQL Query Manager, interactive SQL, embedded SQL, and in any AS/400 HLL. Every message shows up in the job log; you only need to put your job into debug mode. You will find that the help on certain messages sometimes offers hints for improved performance.

The fastest way to look at the optimizer messages during run time of a program is:

- Press the System Request Key.
- Press the Enter key.
- From the System Request menu, select Option 3 (Display current job).
- From the Display Job menu, select Option 10 (Display job log, if active or on job queue).
- On the Display Job Log display, press F10 (Display detailed messages).
- Now page back and locate the optimizer messages.

The display you get after the described procedure may look like this:

```

Display All Messages
Job . . . : DSP010004      User . . . : QPGMR          Number . . . : 002103
5 > opnqryf file((user1/templ))
           qrysl('workdept *eq 'E11'' *and educlvl *GT 17')
All access paths were considered for file TEMPL.
Arrival sequence was used for file TEMPL.

Press Enter to continue.

F3=Exit  F5=Refresh  F12=Cancel  F17=Top  F18=Bottom

```

If you need more information about what the optimizer did, for example:

- Why was arrival sequence used
- Why was an index used
- Why was a temporary index created
- What keys were used to create the temporary index
- What order the files were joined in
- What indexes did the join files use

analyze the messages by pressing the Help key on the message for which you want more information about what happened. If you positioned the cursor on the first message on the previous example display, you may get a display like this:

```

Additional Message Information
Message ID . . . . . : CPI432C          Severity . . . . . : 00
Message type . . . . . : INFO
Date sent . . . . . : 07/08/91          Time sent . . . . . : 09:11:09
From program . . . . . : QQQIMPLE       Instruction . . . . . : 0000
To program . . . . . : QQQIMPLE       Instruction . . . . . : 0000

Message . . . . . : All access paths were considered for file TEMPL.
Cause . . . . . : The OS/400 Query optimizer considered all access paths
built over member TEMPL of file TEMPL in library USER1.
The list below shows the access paths considered. If file TEMPL in
library USER1 is a logical file then the access paths specified are
actually built over member TEMPL of physical file TEMPL in library USER1.
Following each access path name in the list is a reason code which
explains why the access path was not used. A reason code of 0 indicates
that the access path was used to implement the query.
USER1/TEMPLIX1 4, USER1/TEMPLIX2 5, USER1/TEMPLIX3 4.
The reason codes and their meanings follow:
1 - Access path was not in a valid state. The system invalidated the
access path.
Press Enter to continue.                               More...

F3=Exit          F12=Cancel

```

The following is a list of all reason codes the optimizer selects from:

1. Access path was not in a valid state. The system invalidated the access path.
2. Access path was not in a valid state. The user requested that the access path be rebuilt.
3. Access path is a temporary access path (resides in library QTEMP) and was not specified as the file to be queried.
4. The cost to use this access path, as determined by the optimizer, was higher than the cost associated with the chosen access method.

5. The keys of the access path did not match the fields specified for the ordering/grouping criteria.
6. The keys of the access path did not match the fields specified for the join criteria.
7. Use of this access path would not minimize delays when reading records from the file. The user requested to minimize delays when reading records from the file.
8. The access path cannot be used for a secondary file of the join query because it contains static select/omit selection criteria. The join-type of the query does not allow the use of select/omit access paths for secondary files.
9. File xxxx contains record ID selection. The join-type of the query forces a temporary access path to be built to process the record ID selection.
10. The user specified ignore decimal errors on the query. This disallows the use of permanent access paths.
11. The access path contains static select/omit selection criteria that is not compatible with the selection in the query.
12. The access path contains static select/omit selection criteria whose compatibility with the selection in the query could not be determined. Either the select/omit criteria or the query selection became too complex during compatibility processing.
13. The access path contains one or more keys that may be changed by the query during an insert or update.
14. The access path is being deleted or created in an uncommitted unit of work in another process.
15. The keys of the access path match the fields specified for the ordering/grouping criteria. However, the sequence table associated with the access path does not match the sequence table associated with the query.
16. The keys of the access path match the fields specified for the join criteria. However, the sequence table associated with the access path does not match the sequence table associated with the query.
17. The keys of the access path do not match the fields specified for the selection criteria, and the cost to use this access path, as determined by the optimizer, is higher than the cost associated with the chosen access method.
18. The keys of the access path match the fields specified for the selection criteria. However, the sequence table associated with the access path does not match the sequence table associated with the query, and the cost to use this access path, as determined by the optimizer, is higher than the cost associated with the chosen access method.

You can evaluate the performance of your OPNQRYP command using the informational messages put in the job log by the database manager. The database manager may send any of the following messages when appropriate. The ampersand variables (&1, &X) are replacement variables that contain either an object name or another substitution value when the message appears in the job log.

**CPI4321** Access path built for file &1.

**CPI4322** Access path built from keyed file &1.

- CPI4323** The query access plan has been rebuilt.
- CPI4324** Temporary file built for file &1.
- CPI4325** Temporary result file built for query.
- CPI4326** File &1 processed in join position &X.
- CPI4327** File &1 processed in join position 1.
- CPI4328** Access path &4 was used by query.
- CPI4329** Arrival sequence access was used for file &1.
- CPI432A** Query optimizer timed out.
- CPI432B** Subselects processed as join query.
- CPI432C** All access paths were considered for file &1.
- CPI432D** Additional access path reason codes were used.
- CPI432E** Selection fields mapped to different attributes.

The above messages refer not only to OPNQRYF, but also to the SQL/400 or Query/400 programs.

For a more detailed description of the messages and possible user actions, see *Systems Application Architecture\* Structured Query Language/400 Programmer's Guide*, SC41-9609.

---

## Miscellaneous Tips and Techniques

The following is a list of the most important query performance tips:

1. Build the indexes that can be used by queries:
  - Build the indexes to be consistent with the ordering and selection criteria. Use the optimizer messages to help determine the keys needed for consistency. (For more information see "Optimizer Messages" on page D-14.)
  - Avoid putting commonly updated fields in the index.
  - Avoid unnecessary indexes. The optimizer may time out before reaching a good candidate index. (For more information see "Avoiding Too Many Indexes" on page D-18.)
2. Specify the ordering criteria on the primary key of the index to encourage index use when arrival sequence is selected.
3. Include the selection fields in the ordering or grouping criteria. (For more information see "Include Selection fields in Ordering Criteria" on page D-18.)
4. When ordering over a field, try using ALWCPYDTA(\*OPTIMIZE). (For more information see "ORDER BY and ALWCPYDTA" on page D-19.)
5. For %WLDCRD predicate optimization, avoid using the wildcard in the first position. (For more information see "Index Usage with the %WLDCRD Function" on page D-20.)
6. For join optimizations:
  - Avoid joining two files without a JFLD or QRYSLT clause.
  - Create an index on each secondary file.

- Make sure the fields used for joining files match in field length and data type.
- Use redundant predicates for the join fields.
- Allow the primary file to be the file with the fewest number of selected records (for example, no order by or group by on a large file).
- When you specify an ordering on more than 1 file, try using ALWCPYDTA(\*OPTIMIZE).

(For more information about join see “Join Optimization” on page D-20.)

7. Avoid numeric conversion. (For more information see “Avoid Numeric Conversion” on page D-22.)
8. Avoid string truncation by ensuring that literals (constants) are of the same length or shorter than fields being compared to. (For more information see “Avoid String Truncation” on page D-23.)
9. Avoid arithmetic expressions in the selection predicates. (For more information see “Avoid Arithmetic Expressions” on page D-23.)
10. If arrival sequence is used often for queries, use RGZPFM or REUSEDLT(\*YES) to remove deleted records.
11. Use QRYSLT rather than GRPSLT if possible.
12. Avoid mismatching selection and ordering fields.

## Avoiding Too Many Indexes

The available indexes for a query are examined in LIFO (last in, first out) order. If many indexes are available, the optimizer may time out before considering all the indexes. If you are running in debug mode, you will find in the job log the informational message, CPI432A: Query Optimizer timed out for file.

If a timeout occurs, you may delete and re-create an index to put it at the top of the list or use the OPTALLAP option. This option prevents the optimizer from timing out so that all the indexes are considered.

## Include Selection fields in Ordering Criteria

In some situations, you can improve the performance of a query that specifies ordering and selects records based on a *specific* field.

In the following examples that return the same results, note that the second query would be better, because it has the field being selected on as the first key field. This allows *key positioning* to be done instead of just key selection. These are data access methods and are explained in “Access Method” on page D-5.



Instead of:

```
OPNQRYF FILE((TEMPL))
        QRYSLT('EDUCLVL *EQ 16')
        KEYFLD((LASTNAME))
```

Add first key field:

```
OPNQRYF FILE((TEMPL))
        QRYSLT('EDUCLVL *EQ 16')
        KEYFLD((EDUCLVL) (LASTNAME))
```

If no permanent index is available, a temporary index is built, and in the first example, key selection occurs. In the second example, key positioning is done, which is usually much faster.

## ORDER BY and ALWCPYDTA

For most queries that include ordering criteria, the optimizer requires the use of an index. If an index is required and no existing index meets the ordering criteria, OS/400 database support creates a temporary index for this operation.

In some cases, it may be more efficient to sort the records rather than to build an index. In releases prior to Version 2 Release 2, the optimizer did not consider the use of a sort. Now with ALWCPYDTA(\*OPTIMIZE), you can try to improve performance in all the cases in which ordering on the file is required. Sort is considered for read-only cursors when the following values are specified:

- ALWCPYDTA(\*OPTIMIZE) and COMMIT(\*NO)
- ALWCPYDTA(\*OPTIMIZE) and COMMIT(YES) and the journal is opened with a commit level of \*NONE, \*CHG or \*CS

If you specify ALWCPYDTA(\*OPTIMIZE) and have the proper commit level when running the following query,

```
OPNQRYF FILE((TEMPL))
        QRYSLT('DEPTL *EQ ''B01''')
        KEYFLD((NAME))
```

the database manager may:

1. Use an index (on DEPT) or a data scan to resolve the selection criteria.
2. Select the records which meet the selection criteria.
3. Sort the selected records by the values in NAME.

ALWCPYDTA(\*OPTIMIZE) optimizes the total time required to process the query. However, the time required to receive the first row may be increased because a copy of the data must be made prior to returning the first record of the result file. It is very important that an index exist to satisfy the selection criteria. This helps the optimizer to obtain a better estimate of the number of records to be retrieved and to consequently decide whether to use the sort routine or to create a temporary index. Because the sort reads all the query results, the optimizer normally does not perform a sort if OPTIMIZE(\*FIRSTIO) is specified.

Queries that involve a join operation may take advantage of ALWCPYDTA(\*OPTIMIZE). For example, the following join could run faster than the same query with ALWCPYDTA(\*YES) specified:

```
OPNQRYF FILE((FILE1) (FILE2))
        FORMAT(FILE12)
        QRYSLT('FILE1/FLDB *EQ 99 *AND
              FILE2/FLDY *GE 10')
        JFLD((FILE1/FLDA FILE2/FLDX))
        KEYFLD((FLDA) (FLDY))
```

## Index Usage with the %WLDCRD Function

An INDEX will *not* be used when the string in the %WLDCRD function starts with a wildcard character of "%" or "\_". If the string does not start with a wildcard character (for example, "AA%"), the optimizer treats the first byte of the string as a separate selection predicate (for example, "A") and optimizes accordingly, possibly choosing to use an index.

## Join Optimization

The following is a list of join performance tips to help you specify more efficient join operations:

1. Specify join predicates to prevent *all* of the records from one file from being joined to *every* record in the other file:

```
OPNQRYF FILE((FILE1) (FILE2) (FILE3))
        FORMAT(FILE123)
        JFLD((FILE1/FLDA FILE2/FLDA)
              (FILE2/FLDA FILE3/FLDA))
```

In this example, two join predicates are specified.

Each secondary file should have at least one join predicate that references one of its fields as a 'join-to' field.

In the above example, the secondary files, FILE2 and FILE3, both have join predicates that reference FLDA as a join-to field.

2. Create an index over each secondary file

An index is required over each 'join-to' field. If one does not exist, then one is built while the query is processed, which could take a considerable amount of time.

Assuming this is an inner join (JDFTVAL(\*NO)), the optimizer may decide to switch the order of the files specified to a more optimized join order.

To ensure that existing indexes are used for the join operation, create an index over the 'join-from' field and another index over the 'join-to' field of each join predicate. This covers the case where the optimizer could switch the join order.

```
OPNQRYF FILE((TEMP) (TDEPT))
        FORMAT(TEMP)
        QRYSLT('TEMP/WORKBLDG *EQ ''015''')
        JFLD((TEMP/DEPT TDEPT/DEPTNO))
```

The join predicate is TEMP/DEPT = TDEPT/DEPTNO.

In the previous example, the optimizer could perform:

- a. The join from file TEMP to file TDEPT, in which case, an index on TDEPT/DEPTNO, the 'join-to' field, is required.
- b. The join from file TDEPT to file TEMP, in which case, an index on TEMP/DEPT, the 'join-to' field, is required.

Creating indexes on both possible 'join-to' fields, TEMP/DEPT and TDEPT/DEPTNO, covers the case where the optimizer could change the order of joining the files.

3. Specify as many record selection conditions on each file as possible.
4. Specify redundant selection or join predicates, if possible.

Use redundant predicates in a join, if possible, because this allows the optimizer to select the best join order, depending on the estimated number of returned records. Refer to “Key range estimates” in section “Implementation Cost Estimation” on page D-10.

If you are using two files and at the same time qualifying your search condition with a specific value, you should provide a redundant search condition. This helps the optimizer to choose the best way to do the join.

Instead of:

```
OPNQRYF FILE((TEMP) (TDEPT))
        FORMAT(BOTH)
        QRYSLT('TEMP/WORKDEPT *EQ 'E11''')
        JFLD((TEMP/WORKDEPT TDEPT/DEPTNO))
```

Specify:

```
OPNQRYF FILE((TEMP) (TDEPT))
        FORMAT(BOTH)
        QRYSLT('TEMP/WORKDEPT *EQ 'E11'' *AND
        TDEPT/DEPTNO *EQ 'E11''') Redundant predicate
        JFLD((TEMP/WORKDEPT TDEPT/DEPTNO))
```

If you do not duplicate the search condition as shown, then you should put the search condition on the file that selects the fewest records for the predicate.

If the two files being joined have a one-too-many relationship (one record in one file joins to many records in the other file), put the search conditions on the file with unique records.

The example above assumes that TDEPT/DEPTNO is unique, and, therefore, the optimizer chooses a join order of TDEPT as the primary file and TEMP as the secondary join file. In this case, you should build an index over the secondary file (TEMP) with WORKDEPT as the primary key.

5. Make the primary file of the join the file with the fewest number of selected records.

If the specified ordering refers to fields from one file only, that file becomes the primary file. You can force the optimizer to use a file containing the smallest number of selected records as the primary file by specifying a KEYFLD clause that references one of that file's fields.

If the specified ordering refers to fields from one file and that file has a larger number of selected records than its secondary file, consider performance tip 7.

6. Attempt to join the files from smallest to largest, depending on the estimated number of records selected from each file.

The system processes a join of two files with different numbers of selected records most efficiently when the smaller file is joined with the larger file.

The specified join predicates can affect the performance of the join. Always specify the join from the file with the smaller number of selected records to the file with the larger number of selected records, if possible.

7. A temporary result, either forced or allowed through ALWCPYDTA(\*OPTIMIZE), could result in faster join performance. When a temporary result or the sort routine is used, the optimizer is free to reorder the files in the most efficient manner.

If the following conditions apply:

- A file with a larger number of selected records is being joined to a file with a smaller number of selected records.
- The larger file has ordering specified over it that forces it to be the primary file.

then either add a field from the smaller file to the specified ordering, or specify ALWCPYDTA(\*OPTIMIZE).

This allows the use of a temporary file and permits the optimizer to join the files in the most efficient join order.

## Avoid Numeric Conversion

If the file your query is accessing contains numeric fields, you should avoid numeric conversions. As a general guideline, you should always use the same data type for fields and literals used in a comparison. If the data type of the literal has greater precision than the data type of the field, the optimizer will *not* use an index created on that field. To avoid problems for fields and literals being compared, use the:

- Same data type
- Same scale, if applicable
- Same precision, if applicable

In the following example, the data type for the EDUCLVL field is INTEGER. If an index was created on that field, then the optimizer does not use this index in the first OPNQRYF. This is because the precision of the literal is greater than the precision of the field. In the second OPNQRYF, the optimizer considers using the index, because the precisions are equal.

Example where EDUCLVL is INTEGER:

Instead of:

```
OPNQRYF FILE((TEMPL))
  QRYSLT('EDUCLVL *GT 17.0') Index NOT used
```

Specify:

```
OPNQRYF FILE((TEMPL))
  QRYSLT('EDUCLVL *GT 17') Index used
```

## Avoid String Truncation

In general, when literals (constants) are compared to character fields, you should use the same length for the literals as the one specified for the character field:

- If the literals are longer than the field length, the optimizer does *not* use an index created on that field.
- If the literals are shorter than the field length, the optimizer pads the literals with blanks. In the latter case, the optimizer tries to use any available index on the field.

In the following example, the WORKDEPT field is defined as CHAR(3). Assume that an index has been created on that field. The optimizer does *not* use this index in the first query, because the literal compared to WORKDEPT is four bytes long. In the second query, the literal is specified as three bytes, so the optimizer considers using the index.

Example where WORKDEPT is CHAR(3):

Instead of:

```
OPNQRYF FILE((TEMPL))
  QRYSLT('WORKDEPT *EQ 'E11 ''') Index NOT used
```

Specify:

```
OPNQRYF FILE((TEMPL))
  QRYSLT('WORKDEPT *EQ 'E11''') Index used
```

## Avoid Arithmetic Expressions

You should never have an arithmetic expression as an operand to be compared to a field in a record selection predicate. The optimizer does *not* use an index on a field that is being compared to an arithmetic expression.

Instead of:

```
OPNQRYF FILE((TEMPL))
  QRYSLT('SALARY *GT 15000*1.1') Index NOT used
```

Specify:

```
OPNQRYF FILE((TEMPL))
  QRYSLT('SALARY *GT 16500') Index used
```



---

## Bibliography

The following AS/400 manuals contain information you may need. The manuals are listed with their full title and base order number. When these manuals are referred to in this guide, the short title listed is used.

- *Basic Backup and Recovery Guide*, SC41-0036. This manual contains a summary of the information found in the *Advanced Backup and Recovery Guide*, SC41-8079. It provides information about the recovery tools available on the system, such as save and restore operations and power loss recovery. It gives an overview of developing a backup and recovery strategy. It contains procedures and examples for save and restore operations, such as saving and restoring the entire system, saving storage, and restoring licensed internal code. It also contains procedures for disk recovery.

**Short Title:** *Basic Backup and Recovery Guide*.

- *Advanced Backup and Recovery Guide*, SC41-8079. This manual provides information about the recovery tools available on the system, such as save and restore operations, save while active, commitment control, journal management, disk recovery operations and power loss recovery. It also provides guidelines for developing a backup and recovery strategy. It contains procedures for save and restore operations, such as saving and restoring the entire system, saving storage and restoring licensed internal code, and it provides examples of using the save and restore commands. The manual also contains procedures for data and disk recovery, such as using journal management and disk recovering operations, instructions for planning and setting up mirrored protection, and information on uninterruptible power supply. Part 8 contains the appendices for SRC codes, example Disaster Recovery Plan and the IPL process.

**Short Title:** *Advanced Backup and Recovery Guide*.

- *Data Description Specifications Reference*, SC41-9620. This manual provides the application programmer with detailed descriptions of the entries and keywords needed to describe database files (both logical and physical) and certain device files (for displays, printers, and intersystem communications function (ICF)) external to the user's programs.

**Short Title:** *DDS Reference*.

- *Data Management Guide*, SC41-9658. This guide provides the application programmer with information about using files in application programs. Included are topics on the Copy File (CPYF) command and the override commands.

**Short Title:** *Data Management Guide*.

- *Distributed Data Management Guide*, SC41-9600. This guide provides the application programmer with information about remote file processing. It describes how to define a remote file to OS/400 distributed data management (DDM), how to create a DDM file, what file utilities are supported through DDM, and the requirements of OS/400 DDM as related to other systems.

**Short Title:** *DDM Guide*.

- *National Language Support Planning Guide*, GC41-9877. This guide provides the data processing manager, system operator and manager, application programmer, end user, and system engineer with information about understanding and using the national language support function on the AS/400 system. It prepares the user for planning, installing, configuring, and using the AS/400 national language support (NLS) and multilingual system. It also provides an explanation of database management of multilingual data and application considerations for a multilingual system.

**Short Title:** *National Language Support Planning Guide*.

- *Office Services Concepts and Programmer's Guide*, SC41-9758. This guide provides the application programmer with information about writing applications that use OfficeVision/400 functions. It also includes an overview of directory services, document distribution services, document library services, document and folder save and restore and storage management, security services, and word processing services.

**Short Title:** *Office Services Concepts and Programmer's Guide*.

- *Programming: Control Language Programmer's Guide*, SC41-8077. This guide provides the application programmer and programmer with a wide-ranging discussion of AS/400 programming topics, including a general discussion of objects and libraries, CL programming, controlling flow and communicating between programs, working with objects in CL programs, and creating CL programs.

**Short Title:** *CL Programmer's Guide*.

- *Programming: Control Language Reference*, SC41-0030. This set of manuals provides the application programmer and system programmer with detailed information about all AS/400 control language (CL) and its OS/400 commands. All the non-AS/400 CL commands associated with other AS/400 licensed programs, including all the various

languages and utilities are now described in other manuals that support those licensed programs.

**Short Title:** *CL Reference.*

- *Programming: Reference Summary*, SX41-0028. This manual is a quick-reference of various types of summary programming information relating to OS/400 but also to RPG, SEU and SDA. Included are summaries of OS/400 object types, IBM-supplied objects, CL command list, CL command matrix, DDS keywords and monitorable error messages.

**Short Title:** *Programming Reference Summary.*

- *Programming: Work Management Guide*, SC41-8078. This guide provides the programmer with information about how to create and change a work management environment. It also includes a description of tuning the system, collecting performance data including information on record formats and contents of the data being collected, working with system values to control or change the overall operation of the system, and a description of how to gather data to determine who is using the system and what resources are being used.

**Short Title:** *Work Management Guide.*

- *Query/400 User's Guide*, SC41-9614. This guide provides the administrative secretary, business professional, or programmer with information about using AS/400 Query to get data from any database file. It describes how to sign on to Query, and how to define and run queries to create reports containing the selected data.

**Short Title:** *Query/400 User's Guide.*

- *Basic Security Guide*, SC41-0047. This guide explains why security is necessary, defines major

concepts, and provides information on planning, implementing, and monitoring basic security on the AS/400 system.

**Short Title:** *Basic Security Guide.*

- *Security Reference*, SC41-8083. This manual 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.

**Short Title:** *Security Reference.*

- *Systems Application Architecture\* Structured Query Language/400 Programmer's Guide*, SC41-9609. This guide provides the application programmer, programmer, or database administrator with an overview of how to design, write, run, and test SQL/400 statements. It also describes interactive Structured Query Language (SQL).

**Short Title:** *SQL/400\* Programmer's Guide.*

- *Systems Application Architecture\* Structured Query Language/400 Reference*, SC41-9608. This manual provides the application programmer, programmer, or database administrator with detailed information about Structured Query Language/400 statements and their parameters.

**Short Title:** *SQL/400\* Reference.*

- *Utilities: Interactive Data Definition Utility User's Guide*, SC41-9657. This guide provides the administrative secretary, business professional, or programmer with information about using OS/400 interactive data definition utility (IDDU) to describe data dictionaries, files, and records to the system.

**Short Title:** *IDDU User's Guide.*



---

# Index

## Special Characters

\*CT (contains) function and zero length literal 6-9

\*NONE DDS function 3-19, 3-21

## A

**Absolute Value (ABSVAL) keyword** 1-16, 1-24

**ABSVAL (Absolute Value) keyword** 1-16, 1-24

### access method

dynamic processing D-5

index-from-index D-8

key positioning D-7

key selection D-6

### access path

arrival sequence

describing 1-15

reading database records 7-2

attribute 2-3

creating 1-15

definition D-4

describing

overview 1-4

describing logical files 1-15, 3-7

describing physical files 1-15

implicit 3-14

journaling 13-8

keeping current 1-27

keyed sequence

definition 1-16

ignoring 5-4

reading database records 7-3

maximum size A-1

rebuilding

actual time 13-6

controlling 13-6

how to avoid 13-9

reducing time 13-7

recovering

by system 13-5

if the system fails 1-29

restoring 13-5

saving 13-5

select/omit 3-12

sharing 3-13, 3-14, 7-4

specifying

delayed maintenance 1-27

immediate maintenance 1-27

rebuild maintenance 1-27

using

existing specifications 1-23

floating point fields 1-24

access path (*continued*)

writing to auxiliary storage 1-26

**Access Path (ACCPH) parameter** 5-4, 6-2

access plan D-12

**ACCPH (Access Path) parameter** 5-4, 6-2

add authority 4-2

**Add Logical File Member (ADDLFM) command**

DTAMBRS parameter 1-22, 3-21

selecting data members 3-21

using 10-1

**Add Physical File Member (ADDPFM)**

**command** 10-1

adding

logical file member 1-22, 3-21

physical file member 10-1

**ADDLFM (Add Logical File Member) command**

DTAMBRS parameter 1-22, 3-21

using 10-1

**ADDPFM (Add Physical File Member)**

**command** 10-1

**ALIAS (Alternative Name) keyword** 1-9

**ALLOCATE (Allocate) parameter** 2-2

allocating

storage, method 2-2

**Allow Copy Data (ALWCPYDTA) parameter**

ORDER BY field D-19

sort routine D-19

**Allow Delete (ALWDLT) parameter** 2-4, 4-3

**Allow Null (ALWNULL) keyword** 1-9

**Allow Update (ALWUPD) parameter** 2-4, 4-3

alternative collating sequence

arranging key fields 1-16

arranging key fields with SRTSEQ 1-17

**Alternative Name (ALIAS) keyword** 1-9

**ALWCPYDTA (Allow Copy Data) parameter**

ORDER BY D-19

sort routine D-19

**ALWDLT (Allow Delete) parameter** 2-4, 4-3

**ALWNULL (Allow Null) keyword** 1-9

**ALWUPD (Allow Update) parameter** 2-4, 4-3

arithmetic operations using OPNQRYF command

date 6-42

time 6-44

timestamp 6-44

arrival sequence access path

describing 1-15

reading database records 7-2

ascending sequence

arranging key fields 1-18

attribute

database file and member 1-25

source file 14-2

**attribute** (*continued*)

- specifying physical file and member 2-1

**attributes**

- database file and member 1-25

**AUT (Authority) parameter** 1-30, 4-3

**authority**

- add 4-2
- data 4-2
- deleting 4-2
- file and data 4-1
- object 4-1
- public
  - definition 4-3
  - specifying 1-30
- read 4-2
- specifying 4-1
- update 4-2

**Authority (AUT) parameter** 1-30, 4-3

**auxiliary storage**

- writing access paths to
  - frequency 1-26
  - method 5-6
- writing data to
  - frequency 1-26
  - method 5-6

## B

**bibliography** H-1

**blocked input/output**

- See also* sequential-only processing
- improving performance with 5-16

**both fields** 3-3

**bracketed-DBCS data** B-1

## C

**capability**

- database file 4-3
- physical file 2-4

**CCSID (Coded Character Set Identifier)**

**parameter** 1-30

**Change Logical File Member (CHGLFM)**

**command** 10-2

**Change Physical File Member (CHGPFM)**

**command** 10-2

**changing**

- logical file member 10-2
- physical file member 10-2

**Check Expiration Date (EXPCHK) parameter** 5-6

**Check Record Locks (CHKRCDLCK) command** 5-7

**CHGLFM (Change Logical File Member)**

**command** 10-2

**CHGPFM (Change Physical File Member)**

**command** 10-2

**CHKRCDLCK (Check Record Locks) command** 5-7

**Clear Physical File Member (CLRPFM)**

**command** 10-3

**clearing**

- data from physical file members 10-3

**CLOF (Close File) command** 8-1

**Close File (CLOF) command** 8-1

**closing**

- file 8-1

**CLRPFM (Clear Physical File Member)**

**command** 10-3

**CMP (Comparison) keyword** 3-9, 3-13

**coded character set identifier (CCSID)** 1-30

**COLHDG (Column Heading) keyword** 1-9

**collection**

- See* database

**Column Heading (COLHDG) keyword** 1-9

**command**

- database processing options on 5-19
- using output files, example 12-4
- writing output directly to a database file 12-4

**command, CL**

Add Logical File Member (ADDLFM)

- DTAMBRS parameter 1-22, 3-21

- using 10-1

Add Physical File Member (ADDPFM) 10-1

ADDLFM (Add Logical File Member)

- DTAMBRS parameter 1-22, 3-21

- using 10-1

ADDPFM (Add Physical File Member) 10-1

Change Logical File Member (CHGLFM) 10-2

Change Physical File Member (CHGPFM) 10-2

Check Record Locks (CHKRCDLCK) 5-7

CHGLFM (Change Logical File Member) 10-2

CHGPFM (Change Physical File Member) 10-2

CHKRCDLCK (Check Record Locks) 5-7

Clear Physical File Member (CLRPFM) 10-3

CLOF (Close File) 8-1

Close File (CLOF) 8-1

CLRPFM (Clear Physical File Member) 10-3

Copy File (CPYF)

- adding members 10-1

- copying to and from files 14-5

- processing keyed sequence files 1-15

- writing data to and from source file

- members 14-4

Copy from Query File (CPYFRMQRYP) 6-54

Copy Source File (CPYSRCF) 14-4

CPYF (Copy File)

- adding members 10-1

- copying to and from files 14-5

- processing keyed sequence files 1-15

- writing data to and from source file

- members 14-4

CPYFRMQRYP (Copy from Query File) 6-54

CPYSRCF (Copy Source File) 14-4

**command, CL (continued)**

- Create Class (CRTCLS) 7-5
- Create Logical File (CRTLFL)
  - adding members 10-1
  - creating database files 1-24
  - creating source files 14-1
  - DTAMBRS parameter 1-22, 3-21
  - example 3-16
- Create Physical File (CRTPFL)
  - adding members 10-1
  - creating database files 1-24
  - creating source files 14-1
  - RCDLEN parameter 1-3
  - using, example 2-1
- Create Source Physical File (CRTSRCPF)
  - creating physical files 2-1
  - creating source files 14-1
  - describing data to the system 1-3
- CRTCLS (Create Class) 7-5
- CRTLFL (Create Logical File)
  - adding members 10-1
  - creating database files 1-24
  - creating source files 14-1
  - DTAMBRS parameter 1-22, 3-21
  - example 3-16
- CRTPFL (Create Physical File)
  - adding members 10-1
  - creating database files 1-24
  - creating source files 14-1
  - RCDLEN parameter 1-3
  - using, example 2-1
- CRTSRCPF (Create Source Physical File)
  - creating physical files 2-1
  - creating source files 14-1
  - describing data to the system 1-3
  - RCDLEN parameter 1-3
  - using, example 2-1
- Display Database Relations (DSPDBR) 1-14, 12-2
- Display File Description (DSPFD) 12-5, 14-8
- Display File Field Description (DSPFFD) 3-4, 12-1
- Display Journal (DSPJRN) 12-5, 13-3
- Display Message Descriptions (DSPMSGD) 9-1
- Display Object Description (DSPOBJD) 14-8
- Display Physical File Member (DSPPFM) 1-15, 10-5
- Display Problem (DSPPRB) 12-5
- Display Program References (DSPPGMREF) 12-3
- Display Record Locks (DSPRDLCK) 5-7
- DSPDBR (Display Database Relations) 1-14, 12-2
- DSPFD (Display File Description) 12-5, 14-8
- DSPFFD (Display File Field Description) 3-4, 12-1
- DSPJRN (Display Journal) 12-5, 13-3
- DSPMSGD (Display Message Descriptions) 9-1
- DSPOBJD (Display Object Description) 14-8
- DSPPFM (Display Physical File Member) 1-15, 10-5

**command, CL (continued)**

- DSPPGMREF (Display Program References) 12-3
- DSPPRB (Display Problem) 12-5
- DSPRDLCK (Display Record Locks) 5-7
- Edit Object Authority (EDTOBJAUT) 4-3
- EDTOBJAUT (Edit Object Authority) 4-3
- End Journal Access Path (ENDJRNAP) 13-8
- ENDJRNAP (End Journal Access Path) 13-8
- Grant Object Authority (GRTOBJAUT) 4-3
- GRTOBJAUT (Grant Object Authority) 4-3
- Initialize Physical File Member (INZPFM) 7-12, 10-2
- INZPFM (Initialize Physical File Member) 7-12, 10-2
- Open Database File (OPNDBF) 6-1
- OPNDBF (Open Database File) 6-1
- OPNQRYF (Open Query File) 6-1, 6-3
- Override with Database File (OVRDBF) 1-27, 5-1
- OVRDBF (Override with Database File) 1-27, 5-1
- RCLRSC (Reclaim Resources) 8-1
- Reclaim Resources (RCLRSC) 8-1
- Remove Member (RMVM) 10-2
- Rename Member (RNMM) 10-2
- Reorganize Physical File Member (RGZPFM) 7-12, 10-3
- Retrieve Member Description (RTVMBRD) 12-1
- Revoke Object Authority (RVKOBJAUT) 4-3
- RGZPFM (Reorganize Physical File Member) 7-12, 10-3
- RMVM (Remove Member) 10-2
- RNMM (Rename Member) 10-2
- RTVMBRD (Retrieve Member Description) 12-1
- RVKOBJAUT (Revoke Object Authority) 4-3
- Start Journal Access Path (STRJRNAP) 13-8
- Start Journaling Physical File (STRJRNPF) 5-5
- Start Query (STRQRY) 10-5
- Start SQL/400 (STRSQL) 10-5
- STRJRNAP (Start Journal Access Path) 13-8
- STRJRNPF (Start Journaling Physical File) 5-5
- STRQRY (Start Query) 10-5
- STRSQL (Start SQL/400) 10-5
- COMMIT parameter 5-5, 6-2**
- commitment control 5-5, 13-3**
- comparing DBCS fields B-3, B-5**
- Comparison (CMP) keyword 3-9, 3-13**
- CONCAT (Concatenate) keyword 3-1, 3-4**
- Concatenate (CONCAT) keyword 3-1, 3-4**
- concatenated field 3-4**
- concatenating, DBCS B-2**
- concatenation function with DBCS field B-5**
- constant, DBCS B-1**
- contains (\*CT) function and zero length literal 6-9**
- CONTIG (Contiguous Storage) parameter 2-2**
- Contiguous Storage (CONTIG) parameter 2-2**
- conventions, naming 1-5**
- Copy File (CPYF) command**
  - adding members 10-1
  - copying to and from files 14-5

**Copy File (CPYF) command** (*continued*)  
processing keyed sequence files 1-15  
writing data to and from source file members 14-4

**Copy Source File (CPYSRCF) command** 14-4  
**copying**

file  
adding members 10-1  
copying to and from files 14-5  
processing keyed sequence files 1-15  
writing data to and from source file members 14-4  
query file 6-54  
source file 14-4

**correcting errors** 9-1

**CPYF (Copy File) command**

adding members 10-1  
copying to and from files 14-5  
processing keyed sequence files 1-15  
writing data to and from source file members 14-4

**CPYSRCF (Copy Source File) command** 14-4

**Create Class (CRTCLS) command** 7-5

**Create Logical File (CRTLF) command**

adding members 10-1  
creating database files 1-24  
creating source files 14-1  
DTAMBR parameter 1-22, 3-16  
example 3-16

**Create Physical File (CRTPF) command**

adding members 10-1  
creating database files 1-24  
creating source files 14-1  
RCDLEN parameter 1-3  
using, example 2-1

**Create Source Physical File (CRTSRCPF) command**

creating physical files 2-1  
creating source files 14-1  
describing data to the system 1-3  
RCDLEN parameter 1-3  
using, example 2-1

**creating**

class 7-5  
logical file  
adding members 10-1  
creating database files 1-24  
creating source files 14-1  
DTAMBR parameter 1-22, 3-21  
example 3-16  
physical file  
adding members 10-1  
creating database files 1-24  
creating source files 14-1  
DTAMBR parameter 1-22, 3-21  
example 3-16  
source physical file  
creating physical files 2-1  
creating source files 14-1  
describing data to the system 1-3

**CRTCLS (Create Class) command** 7-5

**CRTLF (Create Logical File) command**

adding members 10-1  
creating database files 1-24  
creating source files 14-1  
DTAMBR parameter 1-22, 3-21  
example 3-16

**CRTPF (Create Physical File) command**

adding members 10-1  
creating database files 1-24  
creating source files 14-1  
RCDLEN parameter 1-3  
using, example 2-1

## D

**data**

authority 4-1, 4-2  
clearing from physical file members 10-3  
copying source file 14-4  
describing 1-3  
dictionary-described 1-2  
frequency of writing to auxiliary storage 1-26  
initializing in a physical file member 10-2  
integrity considerations 3-47, 5-5  
recovery considerations 5-5, 13-2  
reorganizing  
physical file member 10-3  
source file members 14-9  
storing 1-26  
using  
default for missing records from secondary files 3-43  
dictionary for field reference 1-13  
example 3-43  
logical files to secure 4-4  
writing to auxiliary storage 13-4

**data description specifications (DDS)**

describing  
database file 1-5  
logical file, example 1-8  
physical file, example 1-5  
using, reasons 1-3

**Data Members (DTAMBR) parameter**

reading order  
logical file members 3-21  
physical file members 1-25

**data space**

definition D-5  
scan D-5

**database**

file attributes 1-25  
member attributes 1-25  
processing options specified on CL commands 5-19  
recovering  
*See also* database file, recovering  
after abnormal system end 13-9

**database** *(continued)*recovering *(continued)*

- data 13-2
- planning 13-1
- restoring 13-1
- saving 13-1
- security 4-1
- using attribute and cross-reference information 12-1

**database file**

- adding members 10-1
- attributes 1-25
- authority types 4-1
- basic operations 7-1
- capabilities 4-3
- changing
  - attributes 11-1
  - descriptions 11-1
- closing
  - methods 8-1
  - sequential-only processing 5-19
  - shared in a job 5-11
- common member operations 10-1
- creating
  - methods 1-24
  - using FORMAT parameter 6-37
- describing
  - methods 1-1
  - to the system 1-4
  - using DDS 1-5
- displaying
  - attributes 12-1
  - descriptions of fields in 12-1
  - information 12-1
  - relationships between 12-2
  - those used by programs 12-3
- estimating size A-1
- grouping data from records 6-32
- handling errors in a program 9-1
- joining without DDS 6-23
- locking
  - considerations C-1
  - wait time 5-8
- minimum size A-2
- naming 5-1
- opening
  - commands to use 6-1
  - members 6-1
  - sequential-only processing 5-16
  - shared in a job 5-9
  - shared in an activation group 5-9
- override 1-27, 5-1
- processing options 5-2
- protecting
  - commitment control 5-5
  - journaling 5-5
- recovering
  - See also* database, recovering

**database file** *(continued)*recovering *(continued)*

- after IPL 13-10
- during IPL 13-9
- options 13-11
- saving and restoring 13-1
- setting a position 7-1
- setting up 1-1
- sharing across jobs 5-6
- sharing in a job
  - close 5-11
  - input/output considerations 5-10
  - open 5-9
  - open data path considerations 6-53
  - SHARE parameter 1-29, 5-8
- sharing in an activation group
  - close 5-11
  - input/output considerations 5-10
  - open 5-9
  - SHARE parameter 5-8
- sizes
  - maximum A-1
  - minimum A-2
- specifying
  - system where created 1-30
  - wait time for locked 1-30
- types 1-25
- with different record formats 6-6
- writing the output from a command to 12-4

**database member**

- adding to files 10-1
- attributes 1-25
- managing 10-1
- naming 5-1
- number allowed 1-25
- removing 10-2

**database record**

- adding 7-9
- deleting 7-11
- file attributes 1-25
- reading methods
  - arrival sequence access path 7-2
  - keyed sequence access path 7-3
- updating 7-8

**database recovery**

*See* database, recovering

**date**

- arithmetic using OPNQRYF command 6-42
- comparison using OPNQRYF command 6-40
- duration 6-41

**DBCS (double-byte character set)**

- considerations B-1
- constant B-1
- field
  - comparing B-3, B-5
  - concatenating B-2
  - concatenation function B-5

**DBCS (double-byte character set) (continued)**

field (continued)

data types B-1

mapping B-2

substring B-3

using the concatenation function B-5

wildcard function B-4

**DDM (distributed data management) 6-38****DDS (data description specifications)**

describing

database file 1-5

logical file, example 1-8

physical file, example 1-5

using, reasons 1-3

**Default (DFT) keyword 1-9, 3-4****default filter factors D-11****defining**

fields 6-28

**definition**

access path D-4

arrival sequence access path D-4

data space D-5

default filter factors D-11

index-from-index access method D-8

key positioning access method D-7

key selection access method D-6

keyed sequence D-4

left-most key D-5

miniplan D-12

optimizer D-3

OS/400 query component D-2

primary key D-5

**delaying**

end-of-file processing 5-4

**Deleted Percentage (DLTPCT) parameter 2-3****deleted record**

reusing 5-3

**deleting**

authority 4-2

database record 2-3, 7-11

**deriving new fields from existing fields 3-4****DESCEND (Descend) keyword 1-19****descending sequence**

arranging key fields 1-18

**describing**

access paths

for database files 1-15

for logical files 3-7

overview 1-4

data to the system 1-3

database file

to the system 1-4

with DDS 1-5

logical file

field use 3-3

floating-point fields in 3-6

record format 3-1

**describing (continued)**

logical file (continued)

with DDS, example 1-8

physical files with DDS

example 1-5

record format 1-4

**description**

checking for changes to the record format 1-27

sharing existing record format 1-13

using existing field 1-9

**design guidelines**

OPNQRYF performance D-1

**designing**

additional named fields 3-1

files to reduce access path rebuild time 13-7

**determining**

auxiliary storage pool 13-1

commit

planning 13-1

data sharing requirements 5-6

duplicate key values 5-5

existing record formats 1-9

field-level security requirements 4-1

if multiple record types are needed in files 3-3

journals 13-2

security requirements 4-1

when a source statement was changed 14-9

which source file member was used to create an object 14-8

**device source file**

using 14-4

**DFT (Default) keyword 1-9, 3-4****dictionary-described data**

definition 1-2

**Display Database Relations (DSPDBR)**

command 1-14, 12-2

**Display File Description (DSPFD) command**

output file 12-5

relating source and objects 14-8

**Display File Field Description (DSPFFD)**

command 3-4, 12-1

**Display Journal (DSPJRN) command**

converting journal receiver entries 13-3

output files 12-5

**Display Message Descriptions (DSPMSGD)**

command 9-1

**Display Object Description (DSPOBJD)**

command 14-8

**Display Physical File Member (DSPPFM)**

command 1-15, 10-5

**Display Problem (DSPPRB) command 12-5****Display Program References (DSPPGMREF)**

command 12-3

**Display Record Locks (DSPRCDLCK)**

command 5-7

**displaying**  
 attributes of files 12-1  
 database relations 1-14, 12-2  
 descriptions of fields in a file 12-1  
 errors 9-1  
 file description 12-5, 14-8  
 file field description 3-4, 12-1  
 files used by programs 12-3  
 information about database files 12-1  
 journal 12-5, 13-3  
 message description 9-1  
 object description 14-8  
 physical file member 1-15, 10-5  
 physical file member records 10-5  
 problem 12-5  
 program reference 12-3  
 record lock 5-7  
 relationships between files on the system 12-2  
 system cross-reference files 12-4

**distributed data management (DDM) 6-38**

**divide by zero**  
 handling 6-31

**DLTPCT (Deleted Percentage) parameter 2-3**

**documentation**  
 using source files for 14-9

**double-byte character set (DBCS)**  
 considerations B-1  
 constant B-1  
 field  
   comparing B-3, B-5  
   concatenating B-2  
   concatenation function B-5  
   data types B-1  
   mapping B-2  
   substring B-3  
   using the concatenation function B-5  
   using the wildcard function B-4

**DSPDBR (Display Database Relations) command 1-14, 12-2**

**DSPFD (Display File Description) command**  
 output file 12-5  
 relating source and objects 14-8

**DSPFFD (Display File Field Description) command 3-4, 12-1**

**DSPJRN (Display Journal) command**  
 converting journal receiver entries 13-3  
 output files 12-5

**DSPMSGD (Display Message Descriptions) command 9-1**

**DSPOBJD (Display Object Description) command 14-8**

**DSPPFM (Display Physical File Member) command 1-15, 10-5**

**DSPPGMREF (Display Program References) command 12-3**

**DSPPRB (Display Problem) command 12-5**

**DSPRCDLCK (Display Record Locks) command 5-7**

**DTAMBRS (Data Members) parameter**  
 reading order  
   logical file members 3-21  
   physical file members 1-25  
   specifying order for files or members 1-22

**DUPKEYCHK (Duplicate Key Check) parameter 5-5, 6-2**

**Duplicate Key Check (DUPKEYCHK) parameter 5-5, 6-2**

**duplicate key field**  
 arranging 1-21  
 preventing 1-20

**duplicate key value 5-5**

**duplicate records in a secondary file**  
 reading 3-34

**duration (date, time, and timestamp) 6-41**

**dynamic access path function 6-21**

**dynamic processing access method D-5**

**Dynamic Select (DYNSLT) keyword 3-12**

**dynamic select/omit 3-12**

**DYNSLT (Dynamic Select) keyword 3-12**

## E

**Edit Code (EDTCDE) keyword 1-9**

**Edit Object Authority (EDTOBJAUT) command 4-3**

**Edit Word (EDTWRD) keyword 1-9**

**EDTCDE (Edit Code) keyword 1-9**

**EDTOBJAUT (Edit Object Authority) command 4-3**

**EDTWRD (Edit Word) keyword 1-9**

**End Journal Access Path (ENDJRNAP) command 13-8**

**end-of-file**  
 delaying processing 5-4  
 waiting for more records 7-5

**ENDJRNAP (End Journal Access Path) command 13-8**

**ensuring data integrity 3-47**

**EOF Retry Delay (EOFDLY) parameter 5-4**

**EOFDLY (EOF Retry Delay) parameter 5-4**

**error**  
 correcting 9-1  
 database file  
   handling in programs 9-1  
 displaying 9-1

**estimating**  
 file size A-1

**example**  
 changing  
   attributes of physical files 11-4  
   descriptions of physical files 11-4  
 closing shared files 5-11  
 complex join logical file 3-45

**example** *(continued)*

- defining
  - fields derived from existing field definitions 6-28
- describing
  - fields that never appear in record format 3-37
  - logical files using DDS 1-8
  - physical files with DDS 1-5
- extra record in secondary file 3-30
- grouping data from database file records 6-32
- handling missing records in secondary join files 6-26
- implicit access path sharing 3-14
- joining
  - database files without DDS 6-23
  - physical file to itself 3-42
  - three or more physical files 3-40
  - two physical files 3-23
- matching records in primary and secondary files 3-28
- performance 6-48
- processing
  - final-total only 6-34
  - unique-key 6-27
- random access 3-30
- reading duplicate records in secondary files 3-34
- record missing in secondary file
  - JDFTVAL keyword not specified 3-28
  - JDFTVAL keyword specified 3-29
- running the OPNQRYF command 6-35
- secondary file has multiple matches for record in primary file 3-29
- selecting records
  - using OPNQRYF command 6-9
  - without using DDS 6-9
- specifying
  - keyed sequence access path without using DDS 6-21
- specifying key fields
  - from different files 6-22
  - join logical file 3-39
- summarizing data from database file records 6-32
- using
  - command output file 12-4
  - default data for missing records from secondary files 3-43
  - join fields whose attributes are different 3-36
  - more than one field to join files 3-33

**existing access path**

- using 3-13

**EXPCHK (Check Expiration Date) parameter 5-6**

**EXPDTE (Expiration Date) parameter**

- changing logical file member 10-2
- specifying 2-1, 5-6

**expiration date**

- checking 5-6
- specifying 2-1

**Expiration Date (EXPDTE) parameter**

- changing logical file member 10-2
- specifying 2-1, 5-6

**F**

**FCFO (First-Changed First-Out) keyword 1-21**

**FEOD (Force-End-Of-Data) operation 7-11**

**field**

- arranging keys 1-16, 1-18
- arranging keys with SRTSEQ 1-17
- both 3-3
- changing in a file description, effects of 11-1
- comparing DBCS B-3, B-5
- concatenating 3-4
- considerations for field use 6-52
- data types, DBCS B-1
- definition
  - See field definition
- deriving new from existing fields 3-4
- describing
  - fields that never appear in record format, example 3-37
  - floating-point in logical files 3-6
  - using logical files 3-3
- displaying descriptions in a file 12-1
- input only 3-3
- join 3-48
- join logical file 3-48
- mapping, DBCS B-2
- neither 3-4
- preventing duplicate key 1-20
- renaming 3-6
- specifying
  - key, example 3-39
  - translation tables 3-6
- substring 3-6
- using
  - data dictionary for reference 1-13
  - existing descriptions and reference files 1-9
  - floating point in access paths 1-24
  - logical files to describe 3-3
  - multiple key 1-19

**field definition**

- derived from existing field definitions 6-28
- functions 1-9

**field reference file**

- definition 1-9

**FIFO (First-In First-Out) keyword 1-21**

**file**

- closing database
  - sequential-only processing 5-19
  - shared in a job 5-11
  - shared in an activation group 5-11
- copying
  - adding members 10-1
  - copying to and from files 14-5



## file (continued)

### copying (continued)

processing keyed sequence files 1-15

writing data to and from source file

members 14-4

creating physical 2-1

creating source 14-1

database

See also database file

attributes 1-25

closing 8-1

options for processing 5-2

processing options 5-2

recovering options 13-2

recovery after IPL 13-10

describing database

to the system 1-4

with DDS 1-5

in a job 6-53

logical

creating 3-15

describing record format 3-1

setting up 3-32

naming 1-25

opening 6-1

physical

creating 2-1

specifying attributes 2-1

sharing

database, across jobs 5-6

database, in the same activation group 5-8

database, in the same job 1-29, 5-8

source 1-25

specifying

member 1-30

text 1-30

**FILE (File) parameter** 5-1

**FILE parameter** 1-25

**FILETYPE (File Type) parameter** 1-25

**filter factor, default** D-11

**final total-only processing** 6-34

**First-Changed First-Out (FCFO) keyword** 1-21

**First-In First-Out (FIFO) keyword** 1-21

**floating point field**

use in access paths 1-24

**FMTSLR (Format Selector) parameter** 7-11

**Force Access Path (FRCACPTH) parameter** 1-26, 5-6

**Force-End-Of-Data (FEOD) operation** 7-11

**Force-Write Ratio (FRCRATIO) parameter**

data integrity considerations 5-6

database data recovery 13-4

specifying file and member attributes 1-26

**FORMAT (Format) keyword** 1-14

**FORMAT (Format) parameter**

OPNQRYF (Open Query File) command 6-21

**FORMAT parameter**

creating a file, considerations 6-37

**Format Selector (FMTSLR) parameter** 7-11

**format, record**

logical file, describing 3-1

**FRCACPTH (Force Access Path) parameter** 1-26, 5-6

**FRCRATIO (Force-Write Ratio) parameter** 1-26, 5-6, 13-4

## G

**Grant Object Authority (GRTOBJAUT)**

command 4-3

**graphic-DBCS constant** B-1

**graphic-DBCS data** B-1

**Group Select (GRPSLT) keyword** 6-35

**grouping**

data from database file records 6-32

performance 6-48

**GRPSLT (Group Select) keyword** 6-35

**GRTOBJAUT (Grant Object Authority)**

command 4-3

## H

**high-level language (HLL) program**

writing considerations 6-38

**HLL (high-level language) program**

writing considerations 6-38

## I

**IBM-supplied source file** 14-2

**IDDU (interactive data definition utility)** 1-3

**ignoring**

keyed sequence access path 5-4

record format 5-5

**implementation cost estimation** D-10

**implicit access path sharing** 3-14

**improving**

performance

for sort sequence 6-48

suggestions 3-46

with OPNQRYF command and keyed sequence

access path 6-45

**index**

See also access path

creating

from another index D-8

fields used for keys D-4

number of D-18

optimizing usage D-21

**Inhibit Write (INHVRT) parameter** 5-6

**INHVRT (Inhibit Write) parameter** 5-6

### initial file position

specifying 5-3

### Initialize Physical File Member (INZPFM)

command 7-12, 10-2

### initializing

data in a physical file member 10-2

### input-only field 3-3

### input/output

blocked

See sequential-only processing

sequential-only processing 5-18

sharing files in a job 5-10

sharing files in an activation group 5-10

### interactive data definition utility (IDDU) 1-3

### INZPFM (Initialize Physical File Member)

command 7-12, 10-2

## J

### JDFTVAL (Join Default Values) keyword 3-28

### JDUPSEQ (Join Duplicate Sequence) keyword 3-31

### JFILE (Joined Files) keyword 3-1

### join

optimization D-13

ordering algorithm D-13

### Join Default Values (JDFTVAL) keyword 3-28

### Join Duplicate Sequence (JDUPSEQ) keyword 3-31

### join field

definition 3-26

rules to remember 3-48

### join logical file

complex, example 3-45

considerations 3-23

definition 3-23

example 3-45

field 3-48

matching records, case 3-28

reading 3-27

requirements 3-47

setting up 3-32

specifying select/omit statements 3-40

summary of rules 3-47

### join optimization

forcing temporary result file D-22

index over secondary file D-20

join from smallest to largest D-22

many record selection conditions D-21

performance tips D-20

predicates on WHERE clause D-20

primary join file has fewest number of selected records D-21

### Join Order (JORDER) parameter 6-23

### Joined Files (JFILE) keyword 3-1

### joining

database files without DDS 6-23

performance 6-48

### joining (continued)

physical file to itself, example 3-42

three or more physical files, example 3-40

two physical files 3-23

two physical files, example 3-23

### JORDER (Join Order) parameter 6-23

### journaling

access path 13-8

commitment control 5-5

definition 13-2

management 13-2

physical file 5-5

## K

### keeping

access paths current 1-27

### key field

arranging

ascending sequence 1-16, 1-18

changing order 1-16

changing order with SRTSEQ 1-17

descending sequence 1-16, 1-18

maximum number, length A-1

preventing duplicate 1-20, 1-21

sharing 7-4

specifying from different files 6-22

subset 7-4

using multiple 1-19

### Key Field (KEYFLD) parameter 6-34

### key positioning access method D-7

### key range estimate D-11

### key selection access method D-6

### keyed sequence D-4

### keyed sequence access path

definition 1-16

reading database records 7-3

### KEYFILE (Key File) parameter 10-3

### KEYFLD (Key Field) parameter 6-34

### keyword, DDS

ABSVAL (Absolute Value) 1-16, 1-24

ALIAS (Alternative Name) 1-9

ALWNULL (Allow Null) 1-9

CMP (Comparison) 3-9, 3-13

COLHDG (Column Heading) 1-9

CONCAT (Concatenate) 3-1, 3-4

DESCEND (Descend) 1-19

DFT (Default) 1-9, 3-4

DYNSLT (Dynamic Selection) 3-12

EDTCDE (Edit Code) 1-9

EDTWRD (Edit Word) 1-9

FCFO (First-Changed First-Out) 1-21

FIFO (First-In First-Out) 1-21

FORMAT (Format) 1-14

GRPSLT (Group Select) 6-35

JDFTVAL (Join Default Values) 3-28

**keyword, DDS** *(continued)*

- JDUPSEQ (Join Duplicate Sequence) 3-31
- JFILE (Joined Files) 3-1
- LIFO (Last-In First-Out) 1-21
- PFILE (Physical File) 1-8, 3-1
- RANGE (Range) 3-8
- REF (Reference) 1-9
- REFACCPH (Reference Access Path definition) 1-23, 1-24, 3-7, 3-8
- REFFLD (Referenced Field) 1-9
- RENAME (Rename) 3-1, 3-6
- SIGNED (Signed) 1-24
- SST (Substring) 3-3, 3-4
- TEXT (Text) 1-9
- TRNTBL (Translation Table) 3-3—3-6
- UNIQUE (Unique)
  - example 1-8
  - preventing duplicate key values 1-20
  - using 1-5, 1-8
- UNSIGNED (Unsigned) 1-16, 1-24
- VALUES (Values) 3-8

**L**

- labeled duration** 6-41
- LANGID (Language Identifier) parameter** 1-31
- language identifier (LANGID)**
  - specifying 1-31
- Last-In First-Out (LIFO) keyword** 1-21
- left-most key** D-5
- length, record** 2-3
- Level Check (LVLCHK) parameter** 1-27, 5-6
- LIFO (Last-In First-Out) keyword** 1-21
- lock**
  - member 5-8
  - record
    - ensuring database integrity 5-7
    - releasing 7-8
    - specifying wait time 1-30
  - record format data 5-8
- logical file**
  - adding 1-21, 3-21
  - adding members 10-1
  - Change Logical File Member (CHGLFM)
    - command 10-2
  - changing
    - attributes 11-5
    - descriptions 11-5
  - creating
    - Create Logical File (CRTLFL) command 1-24
    - database files 1-24
    - DTAMBRS parameter 1-22, 3-16
    - example 3-16
    - methods 3-15
    - source files 14-1
    - with DDS 3-15
    - with more than one record format 3-16

**logical file** *(continued)*

- describing
  - access paths 1-15
  - field use 3-3
  - record format 3-1
  - with DDS, example 1-8
- estimating size A-2
- field
  - describing use 3-3
- join
  - defined 3-23
  - setting up 3-32
- omitting records 3-8
- selecting records 3-8
- setting up 3-1
- sharing access path 7-4
- logical file member** 3-20
- LVLCHK (Level Check) parameter** 1-27, 5-6

**M**

- MAINT (Maintenance) parameter** 1-27
- Maintenance (MAINT) parameter** 1-27
- managing**
  - database member 10-1
  - journals 13-2
  - source file 14-8
- MAPFLD (Mapped Field) parameter** 6-24
- Mapped Field (MAPFLD) parameter** 6-24
- maximum database file sizes** A-1
- Maximum Number of Members (MAXMBRS) parameter** 1-25
- MAXMBRS (Maximum Number of Members) parameter** 1-25
- MBR (Member) parameter**
  - opening members 6-2
  - processing data 5-1
  - specifying member names 1-25
- member**
  - adding to files 10-1
  - attributes 1-25
  - changing attributes 10-2
  - lock 5-8
  - logical file 3-20
  - managing 10-1
  - naming 1-25
  - number allowed in a file 1-25
  - operations common to all database files 10-1
  - removing 10-2
  - renaming 10-2
  - retrieving 12-1
  - source 1-25
  - specifying
    - text 1-30
- Member (MBR) parameter**
  - opening members 6-2

**Member (MBR) parameter** *(continued)*

- processing data 5-1
- specifying member names 1-25

**member description**

- retrieving 12-1

**message**

- sent when OPNQRYF is run 6-38

**minimum database file size** **A-2**

**multiple format logical file**

- adding records 3-20, 7-9
- creating 3-16
- DTAMBRS parameter 3-20
- retrieving records 3-18

## N

**naming**

- database file 5-1
- database member 5-1

**naming conventions** 1-5

**national language support** **B-1**

**NBRRCDS (Number Of Records Retrieved At Once) parameter** 5-16

**neither field** 3-4

**Number Of Records Retrieved At Once (NBRRCDS) parameter** 5-16

## O

**object**

- authority types
  - existence 4-1
  - management 4-1
  - operational 4-1
- creating from source statement in a batch job 14-7

**object authority**

- editing 4-3
- granting 4-3
- revoking 4-3

**OfficeVision/400** **D-2**

**omitting records using logical files** 3-8

**Open Database File (OPNDBF) command** 6-1

**Open File Identifier (OPNID) parameter** 6-2

**Open Query File (OPNQRYF) command**

- running, messages sent 6-38
- using
  - copying 6-54
  - date, time, and timestamp arithmetic 6-41
  - date, time, and timestamp comparison 6-40
  - DBCS fields B-4
  - for more than just input 6-39
  - for random processing 6-45
  - results of a query 6-35
  - selecting records, examples 6-9
  - to select/omit records 3-12
  - typical errors 6-55

**Open Scope (OPNSCOPE) parameter** 6-2

**opening**

- database file
  - commands to use 6-1
  - members 6-1
  - sequential-only processing 5-16
  - shared in a job 5-9
  - shared in an activation group 5-9
- query file 6-1, 6-2

**operation**

- basic database file 7-1
- physical file member 10-2

**OPNDBF (Open Database File) command** 6-1

**OPNID (Open File Identifier) parameter** 6-2

**OPNQRYF (Open Query File) command**

- design guidelines D-1
- performance guidelines D-1
- running, messages sent 6-38
- using
  - copying 6-54
  - date, time, and timestamp arithmetic 6-41
  - date, time, and timestamp comparison 6-40
  - DBCS fields B-4
  - for more than just input 6-39
  - for random processing 6-45
  - results of a query 6-35
  - selecting records, examples 6-9
  - to select/omit records 3-12
  - typical errors 6-55

**OPNSCOPE (Open Scope) parameter** 6-2

**optimizer**

- decision-making rules D-12
- definition D-3
- messages D-14
- operation D-10

**optimizer weightings**

- ALLIO D-10
- FIRSTIO D-10
- MINWAIT D-11

**optimizing**

- index usage D-21

**option**

- database file processing 5-2

**OPTION parameter** 5-2, 6-1

**ORDER BY field**

- ALWCPYDTA D-19

**OS/400 query component** **D-2**

**OUTFILE parameter** 12-4

**output file**

- Display File Description (DSPFD) command 12-5
- Display Journal (DSPJRN) command 12-5
- Display Problem (DSPPRB) command 12-5
- for CL commands 12-4

**Override with Database File (OVRDBF)**

- command** 1-27, 5-1

**OVRDBF (Override with Database File)****command** 1-27, 5-1**P****page fault** D-5**parameter**

ACCPATH (Access Path) 5-4, 6-2  
 ALLOCATE (Allocate) 2-2  
 ALWDLT (Allow Delete) 2-4, 4-3  
 ALWUPD (Allow Update) 2-4, 4-3  
 AUT (Authority) 1-30, 4-3  
 CCSID (Coded Character Set Identifier) 1-30  
 COMMIT 5-5, 6-2  
 CONTIG (Contiguous Storage) 2-2  
 DLTPCT (Deleted Percentage) 2-3  
 DTAMBRS (Data Members)  
   selecting 3-21  
   specifying read order 1-22, 1-25  
 DUPKEYCHK (Duplicate Key Check) 5-5, 6-2  
 EOFDLY (EOF Retry Delay) 5-4  
 EXPCHK (Check Expiration Date) 5-6  
 EXPDATE (Expiration Date)  
   changing of physical file member 10-2  
   specifying expiration date 2-1, 5-6  
 FILE 1-25, 5-1  
 FILETYPE (File Type) 1-25  
 FMTSLR (Format Selector) 7-11  
 FORMAT 6-21, 6-37  
 FRCACCPATH (Force Access Path) 1-26, 5-6  
 FRCRATIO (Force-Write Ratio)  
   data integrity considerations 5-6  
   database data recovery 13-4  
   specifying file and member attributes 1-26  
 INHWRT (Inhibit Write) 5-6  
 JORDER (Join Order) 6-23  
 KEYFILE 10-3  
 KEYFLD (Key Field) 6-34  
 LANGID (Language Identifier) 1-31  
 LVLCHK (Level Check) 1-27, 5-6  
 MAINT (Maintenance) 1-27  
 MAPFLD (Mapped Field) 6-24  
 MAXMBRS (Maximum Number of Members) 1-25  
 MBR (Member)  
   opening members 6-2  
   processing data 5-1  
   specifying member names 1-25  
 NBRRCD (Number Of Records Retrieved At  
 Once) 5-16  
 OPNID (Open File Identifier) 6-2  
 OPNSCOPE (Open Scope) 6-2  
 OPTION 5-2, 6-1  
 OUTFILE 12-4  
 POSITION 5-3, 7-1  
 QRYSLT (Query Select) 3-12  
 RCDFMT (Record Format) 1-14

**parameter (continued)**

RCDFMTLCK (Record Format Lock) 5-8  
 RCDLEN (Record Length) 1-3, 2-3  
 RECORDS 10-2  
 RECOVER 1-29  
 REUSEDLT (Reuse Deleted Records) 2-3  
 SEQONLY (Sequential-Only Processing) 5-16, 6-2  
 SHARE  
   changing for logical files 10-2  
   improving performance 1-29, 5-8  
 SIZE 2-2  
 SRCFILE (Source File) 1-25  
 SRCMBR (Source Member) 1-25  
 SRCOPT (Source Update Options) 10-5, 14-5  
 SRCSEQ (Source Sequence Numbering) 14-5  
 SRCTYPE (Source Type)  
   specifying source type of a member 2-4  
 SRTSEQ (Sort Sequence) 1-30  
 SYSTEM 1-30  
 TEXT 1-30, 10-2  
 TYPE 6-2  
 UNIT 1-26  
 WAITFILE 1-30, 5-8  
 WAITRCD (Wait Record) 1-30, 5-7

**path, access**

*See also* access path  
 creating 1-15  
 recovering  
   if the system fails 1-29

**performance**

arithmetic expressions D-23  
 comparisons with other database functions 6-51  
 considerations D-1  
   for sort sequence 6-48  
   general 6-45  
 examples 6-48  
 grouping, joining, and selection 6-48  
 LIKE predicate D-20  
 numeric conversion D-22  
 OPNQRYF D-1  
 string truncation D-23  
 suggestions 3-46

**performance tools D-2****PFILE (Physical File) keyword** 1-8, 3-1**physical file**

attributes 2-1  
 capabilities 2-4  
 changing  
   attributes 11-2  
   descriptions 11-2  
 creating 2-1  
 CRTPF (Create Physical File) command  
   adding members 10-1  
   creating database files 1-24  
   creating source files 14-1  
 RCDLEN parameter 1-3  
 using, example 2-1

**physical file** *(continued)*

- defined 2-1
- describing
  - access paths 1-15
  - with DDS, example 1-5
- estimating size A-2
- joining
  - three or more, example 3-40
  - to itself, example 3-42
  - two, example 3-23
- journaling
  - starting 5-5
- maximum size, members and key fields A-1
- member size 2-2
- members 2-1
- reorganizing data in members 10-3
- setting up 2-1
- start journaling 5-5
- using
  - DDS to describe, example 1-5
  - existing field descriptions 1-9
  - field reference 1-9

**physical file member**

- adding 10-1
- changing 10-2
- clearing data 10-3
- displaying records 10-5
- initializing data 7-12, 10-2
- reorganizing data 7-11, 10-3
- specifying attributes 2-1

**planning**

- database recovery 13-1

**POSITION parameter** 5-3, 7-1**position, setting in a file** 7-1**pre-fetching** D-6**preventing**

- duplicate key value 1-20
- jobs from changing data in the file 5-6

**primary file**

- definition 3-25

**primary key** D-5**processing**

- database file, options 5-2
- DDM files 6-38
- final total-only 6-34
- options 5-2
- options specified on CL commands 5-19
- random (using OPNQRYF command) 6-45
- sequential-only 5-16
- type of, specifying 5-2
- unique-key 6-27

**program**

- displaying the files used by 12-3
- handling database file errors 9-1
- using source files in 14-6

**protecting**

- file
  - commitment control 5-5
  - journaling 5-5

**public authority**

- definition 4-3
- specifying 1-30

**Q****Q & A database** D-2**QDT (Query Definition Template)** D-12**QRYSLT (Query Select) parameter** 3-12**query**

- starting 10-5

**Query Definition Template (QDT)** D-12**query file**

- copying 6-53
- opening 6-2

**Query Select (QRYSLT) parameter** 3-12**R****random access** 3-30**random processing (using OPNQRYF)** 6-45**RANGE (Range) keyword** 3-8**RCD\_FMT (Record Format) parameter** 1-14**RCD\_FMT\_LCK (Record Format Lock) parameter** 5-8**RCDLEN (Record Length) parameter** 1-3, 2-3**RCLRSC (Reclaim Resources) command** 8-1**reading**

- authority 4-2
- database record, methods
  - arrival sequence access path 7-2—7-3
  - keyed sequence access path 7-3—7-4
- duplicate records in secondary files, example 3-34
- join logical file 3-27

**rebuilding**

- access path 13-6

**Reclaim Resources (RCLRSC) command** 8-1**record**

- adding 7-9
- arranging 6-37
- deleting 2-3, 7-11
- displaying in a physical file member 10-5
- length 2-3
- lock
  - integrity 5-7
  - releasing 7-8
- reading
  - database 7-2
  - physical file 7-2
- reusing deleted 5-3
- specifying
  - length 5-4
  - wait time for locked 1-30

**record** (*continued*)  
 updating 7-8

**record format**  
 checking  
   changes to the description (LVLCHK parameter) 1-27  
   if the description changed, considerations 6-53  
 creating a logical file with more than one 3-16  
 data locks 5-8  
 describing  
   example 1-4  
   logical file 3-1  
 description 1-13  
 ignoring 5-5  
 sharing existing 1-13  
 using  
   different 6-6  
   existing 6-4

**Record Format (RCDFMT) parameter** 1-14  
**Record Format Lock (RCDFMTLCK) parameter** 5-8  
**Record Length (RCDLEN) parameter** 1-3, 2-3

**record lock**  
 checking 5-7  
 displaying 5-7

**record selection method**  
*See* access method

**RECORDS parameter** 10-2  
**RECOVER parameter** 1-29

**recovering**  
*See also* recovery  
 access path  
   by the system 13-5  
 after system end 13-9  
 data 13-2  
 database file  
   after IPL 13-10

**recovery**  
 access path  
   if the system fails 1-29  
 database file  
   during IPL 13-9  
   options table 13-11  
 planning 13-1  
 transaction 13-3

**reducing access path rebuild time** 13-7

**REF (Reference) keyword** 1-9  
**REFACCPH (Reference Access Path Definition) keyword** 1-23, 3-7  
**Reference (REF) keyword** 1-9  
**Reference Access Path Definition (REFACCPH) keyword** 1-23, 3-7  
**Referenced Field (REFFLD) keyword** 1-9  
**REFFLD (Referenced Field) keyword** 1-9

**releasing**  
 locked records 7-8

**Remove Member (RMVM) command** 10-2

**removing**  
 members from files 10-2

**RENAME (Rname) keyword** 3-6  
**RENAME keyword** 3-1  
**Rename Member (RNMM) command** 10-2

**renaming**  
 field 3-6  
 member 10-2

**Reorganize Physical File Member (RGZPFM) command** 7-12, 10-3

**reorganizing**  
 data in physical file members 7-12, 10-3  
 source file member data 14-9

**restoring**  
 access path 13-5  
 data using a disk-resident save file 13-1  
 database  
   functions 13-1

**Retrieve Member Description (RTVMBRD) command** 12-1

**retrieving**  
 member description 12-1  
 records in a multiple format file 3-18

**Reuse Deleted Records (REUSEDLT) parameter** 2-3  
**REUSEDLT (Reuse Deleted Records) parameter** 2-3

**Revoke Object Authority (RVKOBJAUT) command** 4-3

**RGZPFM (Reorganize Physical File Member) command** 7-12, 10-3

**RMVM (Remove Member) command** 10-2  
**RNMM (Rename Member) command** 10-2  
**RTVMBRD (Retrieve Member Description) command** 12-1

**run time**  
 considerations 5-1, 6-54  
 summary 5-19  
 support D-2

**RVKOBJAUT (Revoke Object Authority) command** 4-3

## S

**saving**  
 access path 13-5  
 data using a disk-resident save file 13-1  
 database  
   functions 13-1  
 files and related objects 13-1

**secondary file**  
 definition 3-25  
 example 3-29  
 handling missing records in join 6-26  
 using default data for missing records 3-43

**security**  
 database 4-1

- security** (*continued*)
  - specifying authority 1-30, 4-1
- select/omit**
  - access path 3-12
  - dynamic 3-12
- selecting**
  - record
    - using logical files 3-8
    - using OPNQRYF command 6-9
    - without using DDS, example 6-9
- selection**
  - performance 6-48
- SEQONLY (Sequential-Only Processing)**
  - parameter 5-16, 6-2
- sequence access path**
  - arrival 1-15
  - keyed 1-16
- sequential-only processing**
  - close considerations 5-19
  - input/output considerations 5-18
  - open considerations 5-16
  - SEQONLY parameter 5-16, 6-2
- Sequential-Only Processing (SEQONLY)**
  - parameter 5-16, 6-2
- setting position in file** 7-1
- setting up**
  - database file 1-1
  - join logical file 3-32
  - logical file 3-1
  - physical file 2-1
- SEU (source entry utility)** 14-4
- SHARE (Share) parameter**
  - changing for logical files 10-2
  - improving performance 1-29, 5-8
- sharing**
  - access path 7-4
  - file
    - across jobs 5-6
    - in the same activation group 5-8
    - in the same job 1-29, 5-8
    - OPNQRYF command 6-53
  - implicit access path 3-14
  - record format descriptions that exist 1-13
- SIGNED (Signed) keyword** 1-24
- SIZE parameter** 2-2
- sort sequence**
  - performance considerations 6-48
  - specifying 1-30
- Sort Sequence (SRTSEQ) parameter** 1-30
- source entry utility (SEU)** 14-4
- source file**
  - attributes
    - changing 14-8
    - types 14-2
  - concepts 14-1
  - copying data 14-4
- source file** (*continued*)
  - creating
    - commands 14-1
    - object 14-6
    - with DDS 1-25, 14-3
    - without DDS 14-3
  - entering data 14-3
  - maintaining data 14-3
  - managing 14-8
  - sequence numbers used in copies 14-5
  - statements, determining when changed 14-9
  - supplied by IBM 14-2
  - using
    - device 14-4
    - for documentation 14-9
    - in a program 14-6
- Source File (SRCFILE) parameter** 1-25
- source file member**
  - determining which used to create an object 14-8
  - reorganizing data 14-9
- Source Member (SRCMBR) parameter** 1-25
- source physical file**
  - creating
    - RCDLEN parameter 1-3
    - source files 14-1
    - using, example 2-1
- Source Sequence Numbering (SRCSEQ)**
  - parameter 14-5
- source type**
  - specifying 2-4
- Source Type (SRCTYPE) parameter**
  - specifying 2-4
- Source Update Options (SRCOPT) parameter** 10-5, 14-5
- specifications**
  - using existing access path 1-23
- specifying**
  - access path maintenance levels 1-27
  - attributes
    - physical file and member 2-1
  - authority 4-1
  - database
    - file text 1-30
    - member text 1-30
  - delayed maintenance, access path 1-27
  - expiration date of a file 2-1, 5-6
  - file text 1-30
  - how a file is shared 5-6
  - immediate maintenance, access path 1-27
  - initial file position 5-3
  - key field
    - from different files 6-22
    - in join logical files, example 3-39
  - keyed sequence access path without DDS 6-21
  - LANGID (Language Identifier) 1-31
  - language identifier 1-31



**specifying** (*continued*)

- maximum number of members 1-25
- maximum size of a file A-1
- member attributes 2-1
- member text 1-30
- members, physical files 2-1
- physical file and member attributes 2-1
- physical file attributes 2-1
- public authority 1-30
- rebuild maintenance, access path 1-27
- rebuilt access paths 13-6
- record length 2-3, 5-4
- select/omit statements in join logical files 3-40
- sort sequence 1-30
- source type of a member 2-4
- SRTSEQ (Sort Sequence) parameter 1-30
- system where the file is created 1-30
- type of processing 5-2
- wait time for a locked file or record 1-30

**SQL/400 (Structured Query Language/400) 1-3, 10-5**

**SRCFILE (Source File) parameter 1-25**

**SRCMBR (Source Member) parameter 1-25**

**SRCOPT (Source Update Options) parameter 10-5, 14-5**

**SRCSEQ (Source Sequence Numbering) parameter 14-5**

**SRCTYPE (Source Type) parameter specifying 2-4**

**SRTSEQ (Sort Sequence) parameter 1-30**

**SST (Substring) keyword 3-3, 3-4**

**Start Journal Access Path (STRJRNAP) command 13-8**

**Start Journal Physical File (STRJRNPf) command 5-5**

**Start Query (STRQRY) command 10-5**

**Start SQL/400 (STRSQL) command 10-5**

**starting**

- journal access path 13-8
- journaling physical file 5-5
- query 10-5
- SQL/400 program 10-5

**storage**

- allocating 2-2
- specifying location 1-26
- writing
  - access path to auxiliary 1-26
  - data to auxiliary 1-26, 13-4

**STRJRNAP (Start Journal Access Path) command 13-8**

**STRJRNPf (Start Journal Physical File) command 5-5**

**STRQRY (Start Query) command 10-5**

**STRSQL (Start SQL/400) command 10-5**

**Structured Query Language/400 (SQL/400) 1-3, 10-5**

**Substring (SST) keyword 3-4**

**substring field**

- SST (Substring) keyword 3-6
- using 3-6

**substring operation**

- DBCS B-3
- SST (Substring) keyword 3-6
- using 3-6

**summary**

- database
  - file maximums A-1
  - locks C-1
- rules for join logical files 3-47
- run time 5-19

**SYSTEM parameter 1-30**

**T**

**text**

- specifying
  - database file 1-30
  - database member 1-30
  - file 1-30
  - member 1-30

**TEXT (Text) keyword 1-9**

**TEXT (Text) parameter 1-30, 10-2**

**time**

- arithmetic using OPNQRYF command 6-44
- comparison using OPNQRYF command 6-40
- duration 6-41

**timestamp**

- arithmetic using OPNQRYF command 6-44
- comparison using OPNQRYF command 6-41
- duration 6-41

**tips and techniques for OPNQRYF**

**performance D-17**

**transaction recovery 13-3**

**translated fields 3-6**

**Translation Table (TRNTBL) keyword 3-3—3-6**

**TRNTBL (Translation Table) keyword 3-3—3-6**

**TYPE (Type) parameter 6-2**

**U**

**UNIQUE (Unique) keyword**

- example 1-8
- preventing duplicate key values 1-20
- using 1-5, 1-8

**unique-key processing 6-27**

**UNIT parameter 1-26**

**UNSIGNED (Unsigned) keyword 1-16, 1-24**

**updating**

- authority 4-2
- database record 7-8

**using**

- Open Query File (OPNQRYF) command
- DBCS fields B-4

**using** (*continued*)

wildcard function, DBCS B-4

## **V**

**validation** D-12

**VALUES (Values) keyword** 3-8

## **W**

**Wait Record (WAITRCD) parameter** 1-30, 5-7

**wait time** 1-30

**WAITFILE (Maximum File Wait Time)**

**parameter** 1-30, 5-8

**WAITRCD (Wait Record) parameter** 1-30, 5-7

**wildcard function**

definition B-4

using with DBCS fields B-4

**writing**

access paths to auxiliary storage 5-6

data to auxiliary storage 5-6, 13-4

high-level language program 6-38

output from a command directly to a database

file 12-4

## **Z**

**zero length literal and contains (\*CT) function** 6-9

# Customer Satisfaction Feedback

Application System/400  
Database Guide  
Version 2

Publication No. SC41-9659-02

Overall, how would you rate this manual?

	Very Satisfied	Satisfied	Dissatisfied	Very Dissatisfied
Overall satisfaction				

How satisfied are you that the information in this manual is:

Accurate				
Complete				
Easy to find				
Easy to understand				
Well organized				
Applicable to your tasks				

THANK YOU!

Please tell us how we can improve this manual:

---

---

---

---

---

May we contact you to discuss your responses?  Yes  No

Phone: (\_\_\_\_) \_\_\_\_\_ Fax: (\_\_\_\_) \_\_\_\_\_

**To return this form:**

- Mail it
  - Fax it
  - Hand it to your IBM representative.
- United States and Canada: **800+937-3430**  
Other countries: **(+1)+507+253-5192**

Note that IBM may use or distribute the responses to this form without obligation.

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

ATTN DEPT 245  
IBM CORPORATION  
3605 HWY 52 N  
ROCHESTER MN 55901-9986



Fold and Tape

Please do not staple

Fold and Tape

# Customer Satisfaction Feedback

Application System/400  
Database Guide  
Version 2

Publication No. SC41-9659-02

Overall, how would you rate this manual?

	Very Satisfied	Satisfied	Dissatisfied	Very Dissatisfied
Overall satisfaction				

How satisfied are you that the information in this manual is:

Accurate				
Complete				
Easy to find				
Easy to understand				
Well organized				
Applicable to your tasks				

THANK YOU!

Please tell us how we can improve this manual:

---

---

---

---

---

May we contact you to discuss your responses?  Yes  No

Phone: (\_\_\_\_) \_\_\_\_\_ Fax: (\_\_\_\_) \_\_\_\_\_

To return this form:

- Mail it
  - Fax it
  - Hand it to your IBM representative.
- United States and Canada: **800+937-3430**  
Other countries: **(+1)+507+253-5192**

Note that IBM may use or distribute the responses to this form without obligation.

Name

Address

Company or Organization

Phone No.




Cut  
Along


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

ATTN DEPT 245  
IBM CORPORATION  
3605 HWY 52 N  
ROCHESTER MN 55901-9986



Fold and Tape

Please do not staple

Fold and Tape

Cut c  
Along





Program Number: 5738-SS1

Printed in Denmark by Bonde's

SC41-9659-02

