

IBM Informix Optical Subsystem Guide

Version 9.4
March 2003
Part No. CT1T4NA

Note:

Before using this information and the product it supports, read the information in the appendix entitled "Notices."

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties, and any statements provided in this manual should not be interpreted as such.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1996, 2003. All rights reserved.

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

Table of Contents

Introduction

In This Introduction	3
About This Manual	3
Types of Users	3
Software Dependencies	4
Assumptions About Your Locale.	4
Demonstration Databases	4
New Features	5
Documentation Conventions	5
Typographical Conventions	6
Icon Conventions	7
Syntax Conventions	8
Sample-Code Conventions.	12
Additional Documentation	13
Related Reading	15
Compliance with Industry Standards	16
IBM Welcomes Your Comments	16

Chapter 1

An Overview of the Optical Subsystem

In This Chapter	1-3
The Advantages of Optical Media	1-4
Optical Media and TEXT and BYTE Data	1-4
Components of the Optical-Storage Subsystem	1-5
Optical Platter Organization	1-6
Storage of TEXT and BYTE Data	1-7
Storing TEXT and BYTE Data Objects Sequentially	1-7
Storing TEXT and BYTE Data in a Cluster	1-7
Placing TEXT and BYTE Data in Optical Clusters	1-9
Calculating Logical-Log Space for Optical Data	1-11

Support for Multiple Optical Virtual Processors	1-11
Determining the Number of Optical Virtual Processors	1-11
The Optical Subsystem API	1-12
The Memory Cache and Staging-Area BlobSpace	1-14
The API Internal Layer	1-19
The API External Layer	1-21
The Optical Subsystem Message File	1-21

Chapter 2 Installation and Initial Configuration

In This Chapter	2-3
Prerequisites	2-4
Supporting the Shared Library	2-4
Creating the Staging Area	2-5
Naming the Staging Area on the STAGEBLOB Parameter	2-5
Configuring Optical Virtual Processors	2-6
Initializing the Optical Subsystem	2-6
Creating the Staging-Area BlobSpace	2-7
Verifying the Presence of the Optical Subsystem	2-9
Creating Optical Families	2-9
Testing the Connection	2-10
Test One	2-10
Test Two	2-10

Chapter 3 Using the Optical Subsystem

In This Chapter	3-3
Assigning TEXT and BYTE Columns to an Optical Platter	3-4
Reading TEXT and BYTE Data Columns from an Optical Platter	3-5
Clustering TEXT and BYTE Data	3-5
Managing Volumes on the Optical Drives	3-9
Using TEXT and BYTE Data Descriptors	3-12
Locating the Optical Volume Where TEXT and BYTE Data Is Stored	3-14
Locating Columns Stored in Optical Families	3-14
Migrating a Database with TEXT and BYTE Data on an Optical Platter	3-15
Using HPL to Unload and Load Optical Data	3-16
The dbexport and dbimport Utilities	3-16

Chapter 4	SQL for the Optical Subsystem	
	In This Chapter	4-3
	ALTER OPTICAL CLUSTER	4-4
	CREATE OPTICAL CLUSTER.	4-6
	DROP OPTICAL CLUSTER	4-10
	RELEASE	4-12
	RESERVE	4-14
	SET MOUNTING TIMEOUT	4-16
	Function Expressions	4-18
Appendix A	Notices	
	Index	

Introduction

In This Introduction	3
About This Manual.	3
Types of Users	3
Software Dependencies	4
Assumptions About Your Locale.	4
Demonstration Databases	4
New Features.	5
Documentation Conventions	5
Typographical Conventions	6
Icon Conventions	7
Comment Icons	7
Feature, Product, and Platform Icons	7
Syntax Conventions	8
Elements That Can Appear on the Path	9
How to Read a Syntax Diagram.	11
Sample-Code Conventions	12
Additional Documentation	13
Related Reading.	15
Compliance with Industry Standards	16
IBM Welcomes Your Comments	16

In This Introduction

This introduction provides an overview of the information in this manual and describes the conventions it uses.

About This Manual

This manual describes the Optical Subsystem, a configuration of IBM Informix Dynamic Server that supports the storage of TEXT and BYTE data types, also known as *simple large objects*, on optical platters. It describes the Optical Subsystem interface for an optical-storage subsystem and the set of SQL statements that you use exclusively with optical platters.

Types of Users

This manual is for the following users who want to take advantage of write-once-read-many (WORM) optical media for storage of simple large objects:

- Database administrators
- Database server administrators

This manual assumes that you have the following background:

- A working knowledge of your computer, your operating system, and the utilities that your operating system provides
- Some experience working with relational databases or exposure to database concepts

If you have limited experience with relational databases, SQL, or your operating system, refer to *IBM Informix Dynamic Server Getting Started Guide* for a list of supplementary titles.

Software Dependencies

This manual assumes that you are using IBM Informix Dynamic Server.

Assumptions About Your Locale

IBM Informix products can support many languages, cultures, and code sets. All culture-specific information is brought together in a single environment, called a Global Language Support (GLS) locale.

This manual assumes that you use the U.S. 8859-1 English locale as the default locale. The default is **en_us.8859-1** (ISO 8859-1) on UNIX platforms or **en_us.1252** (Microsoft 1252) for Windows environments. This locale supports U.S. English format conventions for dates, times, and currency, and also supports the ISO 8859-1 or Microsoft 1252 code set, which includes the ASCII code set plus many 8-bit characters such as é, è, and ñ.

If you plan to use nondefault characters in your data or your SQL identifiers, or if you want to conform to the nondefault collation rules of character data, you need to specify the appropriate nondefault locale.

For instructions on how to specify a nondefault locale, additional syntax, and other considerations related to GLS locales, see the *IBM Informix GLS User's Guide*.

Demonstration Databases

The DB-Access utility, which is provided with your IBM Informix database server products, includes one or more of the following demonstration databases:

- The **stores_demo** database illustrates a relational schema with information about a fictitious wholesale sporting-goods distributor. Many examples in IBM Informix manuals are based on the **stores_demo** database.
- The **superstores_demo** database illustrates an object-relational schema. The **superstores_demo** database contains examples of extended data types, type and table inheritance, and user-defined routines.

For information about how to create and populate the demonstration databases, see the *IBM Informix DB-Access User's Guide*. For descriptions of the databases and their contents, see the *IBM Informix Guide to SQL: Reference*.

The scripts that you use to install the demonstration databases reside in the **\$INFORMIXDIR/bin** directory on UNIX platforms and in the **%INFORMIXDIR%\bin** directory in Windows environments.

New Features

For a comprehensive list of new features for your database server, see the *Getting Started Guide*. This section lists new features relevant to this manual.

One feature that affects IBM Informix Optical Subsystem is long identifiers. This feature increases the length of all identifiers from 18 characters to 128 characters. The maximum length of user names is increased from 8 to 32 characters. For more information on long identifiers, see the *IBM Informix Guide to SQL: Syntax*.

Documentation Conventions

This section describes the conventions that this manual uses. These conventions make it easier to gather information from this and other volumes in the documentation set.

The following conventions are discussed:

- Typographical conventions
- Icon conventions
- Syntax conventions
- Sample-code conventions

Typographical Conventions

This manual uses the following conventions to introduce new terms, illustrate screen displays, describe command syntax, and so forth.

Convention	Meaning
KEYWORD	All primary elements in a programming language statement (keywords) appear in uppercase letters in a serif font.
<i>italics</i> italics <i>italics</i>	Within text, new terms and emphasized words appear in italics. Within syntax and code examples, variable values that you are to specify appear in italics.
boldface <i>boldface</i>	Names of program entities (such as classes, events, and tables), environment variables, file and pathnames, and interface elements (such as icons, menu items, and buttons) appear in boldface.
<code>monospace</code> <i>monospace</i>	Information that the product displays and information that you enter appear in a monospace typeface.
KEYSTROKE	Keys that you are to press appear in uppercase letters in a sans serif font.
◆	This symbol indicates the end of one or more product- or platform-specific paragraphs.
→	This symbol indicates a menu item. For example, “Choose Tools→Options ” means choose the Options item from the Tools menu.






Tip: When you are instructed to “enter” characters or to “execute” a command, immediately press RETURN after the entry. When you are instructed to “type” the text or to “press” other keys, no RETURN is required.

Icon Conventions

Throughout the documentation, you will find text that is identified by several different types of icons. This section describes these icons.





Comment Icons

Comment icons identify three types of information, as the following table describes. This information always appears in *italics*.

Icon	Label	Description
	<i>Warning:</i>	Identifies paragraphs that contain vital instructions, cautions, or critical information
	<i>Important:</i>	Identifies paragraphs that contain significant information about the feature or operation that is being described
	<i>Tip:</i>	Identifies paragraphs that offer additional details or shortcuts for the functionality that is being described

Feature, Product, and Platform Icons

Feature, product, and platform icons identify paragraphs that contain feature-specific, product-specific, or platform-specific information.


Icon	Description
	Identifies information that is specific to IBM Informix Dynamic Server
	Identifies information that relates to the IBM Informix Global Language Support (GLS) feature
	Identifies information that is specific to UNIX platforms
	Identifies information that is specific to the Windows environment

These icons can apply to a row in a table, one or more paragraphs, or an entire section. If an icon appears next to a section heading, the information that applies to the indicated feature, product, or platform ends at the next heading at the same or higher level. A ♦ symbol indicates the end of the feature-, product-, or platform-specific information that appears within a table or a set of paragraphs within a section.

Syntax Conventions

This section describes conventions for syntax diagrams. Each diagram displays the sequences of required and optional keywords, terms, and symbols that are valid in a given statement or segment, as [Figure 1](#) shows.

Figure 1
Example of a Simple Syntax Diagram











Each syntax diagram begins at the upper-left corner and ends at the upper-right corner with a vertical terminator. Between these points, any path that does not stop or reverse direction describes a possible form of the statement.




Syntax elements in a path represent terms, keywords, symbols, and segments that can appear in your statement. The path always approaches elements from the left and continues to the right, except in the case of separators in loops. For separators in loops, the path approaches counterclockwise. Unless otherwise noted, at least one blank character separates syntax elements.

Elements That Can Appear on the Path

You might encounter one or more of the following elements on a path.

Element	Description
KEYWORD	A word in UPPERCASE letters is a keyword. You must spell the word exactly as shown; however, you can use either uppercase or lowercase letters.
(. , ; @ + * - /)	Punctuation and other nonalphanumeric characters are literal symbols that you must enter exactly as shown.
' '	Single quotes are literal symbols that you must enter as shown.
<i>variable</i>	A word in <i>italics</i> represents a value that you must supply. A table immediately following the diagram explains the value.
<div>ADD Clause p. 1-14</div> <div>ADD Clause</div>	A reference in a box represents a subdiagram. Imagine that the subdiagram is spliced into the main diagram at this point. When a page number is not specified, the subdiagram appears on the same page.
<div>Table Name see Database Object Name in SQLS</div>	A reference to SQLS in this manual refers to the <i>Informix Guide to SQL: Syntax</i> . Imagine that the subdiagram is spliced into the main diagram at this point.
<div>E/C</div>	<p>An icon is a warning that this path is valid only for some products, or only under certain conditions. Characters on the icons indicate what products or conditions support the path.</p> <p>These icons might appear in a syntax diagram:</p> <div><div>DB</div> This path is valid only for DB-Access.</div> <div><div>E/C</div> This path is valid only for IBM Informix ESQL/C.</div> <div><div>EXT</div> This path is valid for external routines.</div>

Element	Description
	This path is valid only if your database or application uses a nondefault GLS locale.
	This path is valid only for the Optical Subsystem.
	This path is valid only if you are using an Informix SPL routine.
	This path is valid for the SQL Editor.
	This path is an Informix extension to ANSI SQL-92 entry-level standard SQL. If you initiate Informix extension checking and include this syntax branch, you receive a warning. If you have set the DBANSIWARN environment variable at compile time, or have used the -ansi compile flag, you receive warnings at compile time. If you have DBANSIWARN set at runtime, or if you compiled with the -ansi flag, warning flags are set in the sqlwarn structure.
	A shaded option is the default action.
	Syntax that is enclosed between a pair of arrows is a subdiagram.
	The vertical line terminates the syntax diagram.


Element	Description
	A branch below the main path indicates an optional path. (Any term on the main path is required, unless a branch can circumvent it.)
	A set of multiple branches indicates that a choice among more than two different paths is available.
	A loop indicates a path that you can repeat. Punctuation along the top of the loop indicates the separator symbol for list items. If no symbol appears, a blank space is the separator.

(3 of 3)

How to Read a Syntax Diagram

Figure 2 shows a syntax diagram that uses most of the path elements that the previous table lists.

Figure 2
Example of a Syntax Diagram



To use this diagram to construct a statement, start at the top left with the keyword `DELETE FROM`. Then follow the diagram to the right, proceeding through the options that you want.

Figure 2 illustrates the following steps:

1. Type `DELETE FROM`.
2. You can delete a table, view, or synonym:
 - Type the table name, view name, or synonym, as you desire.
 - You can type `WHERE` to limit the rows to delete.
 - If you type `WHERE` and you are using DB-Access or the SQL Editor, you must include the Condition clause to specify a condition to delete. To find the syntax for specifying a condition, go to the Condition segment on the specified page.
 - If you are using ESQL/C, you can include either the Condition clause to delete a specific condition or the `CURRENT OF cursor` clause to delete a row from the table.
3. Follow the diagram to the terminator.
Your `DELETE` statement is complete.

Sample-Code Conventions

Examples of SQL code occur throughout this manual. Except where noted, the code is not specific to any single IBM Informix application development tool. If only SQL statements are listed in the example, they are not delimited by semicolons. For instance, you might see the code in the following example:

```
CONNECT TO stores_demo
...

DELETE FROM customer
      WHERE customer_num = 121
...

COMMIT WORK
DISCONNECT CURRENT
```

To use this SQL code for a specific product, you must apply the syntax rules for that product. For example, if you are using DB-Access, you must delimit multiple statements with semicolons. If you are using an SQL API, you must use `EXEC SQL` at the start of each statement and a semicolon (or other appropriate delimiter) at the end of the statement.



Tip: *Ellipsis points in a code example indicate that more code would be added in a full application, but it is not necessary to show it to describe the concept being discussed.*

For detailed directions on using SQL statements for a particular application development tool or SQL API, see the manual for your product.

Additional Documentation

IBM Informix Dynamic Server documentation is provided in a variety of formats:

- **Online manuals.** The Informix OnLine Documentation Web site at <http://www-3.ibm.com/software/data/informix/pubs/library/> contains manuals for your use. This Web site enables you to print chapters or entire books.
- **Online help.** This facility can provide context-sensitive help, an error message reference, language syntax, and more.
- **Documentation notes and release notes.** Documentation notes, which contain additions and corrections to the manuals, and release notes are located in the directory where the product is installed.

Please examine these files because they contain vital information about application and performance issues. The following table describes these files.

UNIX

On UNIX platforms, the following online files appear in the \$INFORMIXDIR/release/en_us/0333 directory.

Online File	Purpose
ids_optical_docnotes_9.40.html	The documentation notes file for your version of this manual describes topics that are not covered in the manual or that were modified since publication.
ids_unix_release_notes_9.40.html ids_unix_release_notes_9.40.txt	The release notes file describes feature differences from earlier versions of IBM Informix products and how these differences might affect current products. This file also contains information about any known problems and their workarounds.
ids_machine_notes_9.40.txt	The machine notes file describes any special actions that you must take to configure and use IBM Informix products on your computer. Machine notes are named for the product described.



Windows

The following items appear in the **Informix** folder. To display this folder, choose **Start→Programs→Informix→ Documentation Notes or Release Notes** from the task bar.

Program Group Item	Description
Documentation Notes	This item includes additions or corrections to manuals with information about features that might not be covered in the manuals or that have been modified since publication.
Release Notes	This item describes feature differences from earlier versions of IBM Informix products and how these differences might affect current products. This file also contains information about any known problems and their workarounds.

Machine notes do not apply to Windows platforms. ♦

- IBM Informix software products provide ASCII files that contain all of the error messages and their corrective actions. For a detailed description of these error messages, refer to *IBM Informix Error Messages* in the IBM Informix Online Documentation site at <http://www-3.ibm.com/software/data/informix/pubs/library/>.

To read the error messages on UNIX, you can use the **finderr** command to display the error messages online. ♦

To read error messages and corrective actions on Windows, use the **Informix Error Message** utility. To display this utility, choose **Start→Programs→Informix** from the task bar. ♦

UNIX

Windows

Related Reading

For a list of publications that provide an introduction to database servers and operating-system platforms, refer to your *Getting Started Guide*.

Compliance with Industry Standards

The American National Standards Institute (ANSI) has established a set of industry standards for SQL. IBM Informix SQL-based products are fully compliant with SQL-92 Entry Level (published as ANSI X3.135-1992), which is identical to ISO 9075:1992. In addition, many features of Informix database servers comply with the SQL-92 Intermediate and Full Level and X/Open SQL CAE (common applications environment) standards.

IBM Welcomes Your Comments

To help us with future versions of our manuals, let us know about any corrections or clarifications that you would find useful. Include the following information:

- The name and version of your manual
- Any comments that you have about the manual
- Your name, address, and phone number

Send electronic mail to us at the following address:

`docinf@us.ibm.com`

This address is reserved for reporting errors and omissions in our documentation. For immediate help with a technical problem, contact Customer Services.

An Overview of the Optical Subsystem

In This Chapter	1-3
The Advantages of Optical Media.	1-4
Optical Media and TEXT and BYTE Data	1-4
Components of the Optical-Storage Subsystem	1-5
Optical Platter Organization.	1-6
Storage of TEXT and BYTE Data	1-7
Storing TEXT and BYTE Data Objects Sequentially	1-7
Storing TEXT and BYTE Data in a Cluster	1-7
Placing TEXT and BYTE Data in Optical Clusters	1-9
TEXT and BYTE Data Columns	1-9
The Cluster-Key Column	1-9
How Optical Clustering Occurs.	1-10
Calculating Logical-Log Space for Optical Data.	1-11
Support for Multiple Optical Virtual Processors	1-11
Determining the Number of Optical Virtual Processors	1-11
The Optical Subsystem API	1-12
The Memory Cache and Staging-Area BlobSpace	1-14
Memory Cache	1-16
The Staging Area	1-16
The STAGEBLOB Parameter	1-17
The onstat -O Option	1-18
The API Internal Layer	1-19
Memory for the Optical Transfer Buffer	1-20
The API External Layer	1-21
The Optical Subsystem Message File	1-21

In This Chapter

The Optical Subsystem is a functionality of Informix database servers that supports the storage of TEXT and BYTE data on an optical platter.

The Optical Subsystem includes an interface for an optical-storage subsystem and supports a set of SQL statements that are exclusively for use with optical platters. These SQL statements support the efficient storage and retrieval of data to and from the optical-storage subsystem.

For more information on TEXT and BYTE data types, refer to the *IBM Informix Guide to SQL: Reference*.



Important: *The Optical Subsystem does not store CLOB and BLOB data types. The optical-storage subsystem uses a file interface to support smart large objects. For more information on how to store CLOB and BLOB data types on optical disc, refer to the “IBM Informix ESQL/C Programmer’s Manual.”*

This chapter describes the purpose and implementation of the Optical Subsystem. The following topics are covered:

- The advantages that optical media offers
- The components of an optical-storage subsystem
- The ways in which you can organize data on optical platters
- The sequential and clustered methods of storing TEXT and BYTE data on optical platters
- The operation of the Optical Subsystem application-programming interface (API), which supports the optical-storage subsystem

The Advantages of Optical Media

Optical media offer the following advantages for storing data over conventional magnetic disks:

- Mass storage capacity (on the order of gigabytes)
- Mountable/unmountable storage units
- Low cost per bit of storage
- Long media life
- High data stability

Because optical media are usually written to and read with a laser, a bit of data is written to a much smaller area on the optical *platter* (disc) than the area that is required to record a bit on a conventional magnetic disk. Optical platters offer a high degree of stability because they are far less susceptible than magnetic disks to corruption by environmental influences such as electromagnetic fields. Furthermore, the possibility of a read/write head crash is extremely remote because the optical disc read/write head does not come as close to the platter as the magnetic disk read/write head.

The Optical Subsystem supports only the write-once-read-many (WORM) type of optical media. When data is written to WORM optical media, a permanent, virtually incorruptible mark is made on the platter. Thereafter, you can read the data an unlimited number of times. When you update the data to a WORM optical platter, it is written to a new location on the platter; the optical-storage subsystem cannot reuse the original location.

Optical Media and TEXT and BYTE Data

In a database environment, the vast storage capacity of optical media make it particularly attractive for storing TEXT and BYTE data such as text documents, source and object programs, and scanned images. TEXT and BYTE data theoretically has no maximum size, although in practice the Optical Subsystem allows a maximum size of about 2 gigabytes per TEXT and BYTE data object. The number of locks available in the system might impose an additional constraint on the size.

The trade-off for the capacity, lower cost, integrity, and security that optical media provide, however, is slower access. In general, accessing optical media is slower than accessing magnetic disks. Typically, an optical drive accesses data at speeds that are 25 to 50 percent slower than a magnetic disk drive.

Components of the Optical-Storage Subsystem

The optical-storage subsystem has the following components:

- The platter
- The jukebox (autochanger) or stand-alone drive
- The shelf
- Management software for the optical-storage subsystem

The optical media are the removable optical platters that contain data.

You can keep optical platters in an automated library called a *jukebox* or *autochanger*. The jukebox is a cabinet that contains one or more optical platter drives, slots that store the optical platters when they are not mounted, and a robotic arm that transfers platters between the slots and the drives. Optical-storage subsystem commands initiate transport of the platters to and from the drives and manage the input/output (I/O) operations within the optical-storage subsystem.

Jukeboxes are not essential to the optical-storage subsystem. A supported optical-storage subsystem can consist of a stand-alone optical drive that holds only one platter or several stand-alone drives.

You store optical platters off-line on *shelves*. When a platter is known to the optical-storage subsystem but is not stored in the jukebox where it can be accessed, the platter is said to be *on the shelf*. The operator is responsible for storage and retrieval of platters to and from the shelf.

The management software for the optical-storage subsystem performs various tasks related to the optical-storage subsystem. It creates optical families and volumes and tracks information about them, chooses which optical drives to use and the optical platters to mount, issues commands to the jukebox, and reads and writes data to the optical platter.

Optical Platter Organization


In a WORM optical-storage subsystem that the Optical Subsystem supports, the optical media are organized into units of storage called the *volume* and the *family*. Each side of the platter is called a volume. An optical family is theoretically an unlimited collection of volumes, although in practice the constraints of the optical-storage subsystem impose a limit on its size. When a volume becomes full, the optical-storage subsystem automatically allocates another volume for the family.

The optical-storage administrator creates an optical family with the management software that the vendor of the optical-storage subsystem provides. The administrator also uses optical-storage subsystem commands to add volumes to a family.

In a manner that is similar to an automated tape library, the optical-storage subsystem tracks the location of each TEXT and BYTE data object on a volume, each volume on a platter, and each platter in a family. Subsystem software schedules the optical drives and mounts platters as needed.

Figure 1-1 shows the relationship between the optical family, the volume, and the optical platter.

Figure 1-1
The Optical Family



Storage of TEXT and BYTE Data

You assign a TEXT or BYTE column to the optical-storage subsystem with the CREATE TABLE statement, which allows you to place the column data in an optical family. Prior to executing the CREATE TABLE statement, the optical-storage administrator must create an optical family name with the management software that the vendor of the optical-storage subsystem provides. You can then use the column-definition portion of the CREATE TABLE statement to assign the column to the optical family. For more information, see [“Assigning TEXT and BYTE Columns to an Optical Platter” on page 3-4](#).

TEXT and BYTE data stored on optical media are not backed up (or archived) by the **ontape**, **onload**, and **onunload** utilities. However, the **dbexport** and **dbimport** utilities can be used to unload and load TEXT and BYTE data stored on optical media.

The Optical Subsystem can store TEXT and BYTE data on an optical volume either sequentially or in a cluster. The following sections describe each of these methods.

Storing TEXT and BYTE Data Objects Sequentially

Sequential storage means that the Optical Subsystem stores TEXT and BYTE data on the current volume of the family in the same sequence that they are inserted. When a volume becomes full, the optical-storage subsystem allocates the next volume in the family; no attempt is made to keep logically related TEXT and BYTE data objects together on the same volume or even on the same platter.

Storing TEXT and BYTE Data in a Cluster

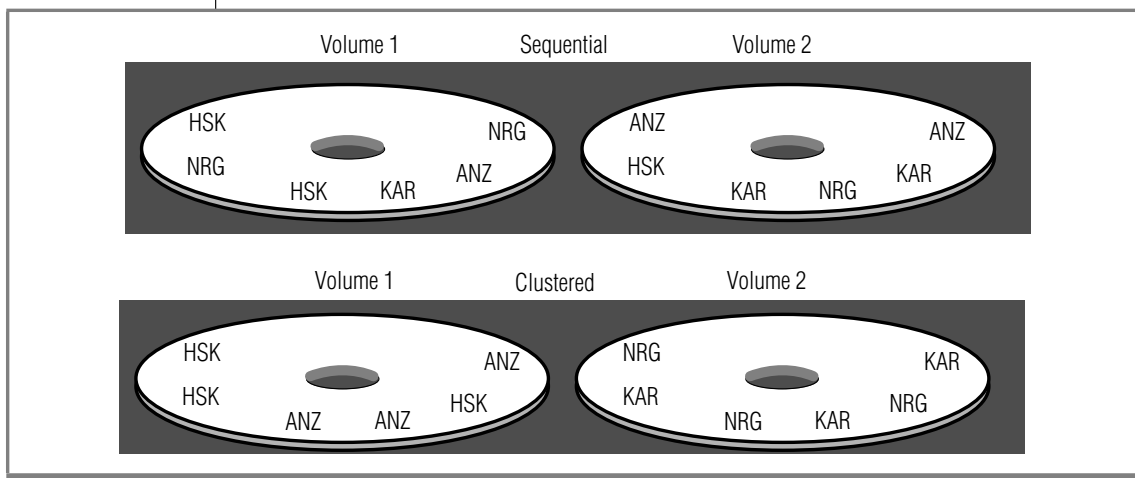
Optical clustering attempts to store logically related TEXT or BYTE data objects together on the same volume. In applications where related TEXT or BYTE data objects are often retrieved together, optical clustering minimizes time-consuming platter exchanges on the drives. This practice improves performance.

Optical clustering does not imply that TEXT or BYTE data objects that share a particular key value are stored next to each other on the optical platter. Instead, each optical cluster represents a reservation of some number of kilobytes (the *clustersize*) on a volume. However, the reserved kilobytes and the TEXT or BYTE data objects that are eventually stored might not be contiguous. The benefit of clustering is that you can find related TEXT or BYTE data objects on the same volume far faster than you can find them on separate volumes or platters.

Figure 1-2 illustrates the difference between sequential and clustered organization. In this example, the optical platters store TEXT or BYTE data objects that are scanned pictures of merchandise. The pictures belong to the **cat_descr** column in the **stores_demo** database. If you cluster the pictures, you must cluster them by reference to some other column, such as the manufacturer's name, as pictures do not have an inherent order themselves.

When stored sequentially, the optical-storage subsystem writes the pictures to the optical platter in the same order that they are inserted in the table. If you cluster them by the manufacturer code (the cluster key), the pictures for a given manufacturer are stored together. The clustered organization reduces platter exchanges on the optical drives.

Figure 1-2
Sequential Storage Versus Clustering by manu_code



Placing TEXT and BYTE Data in Optical Clusters

You create an optical cluster with the CREATE OPTICAL CLUSTER statement. For information on the syntax, refer to “[CREATE OPTICAL CLUSTER](#)” on [page 4-6](#). As part of the syntax to define the cluster, specify the following column lists:

- | | |
|--------------------------------|---|
| <i>column list</i> | is one or more logically related TEXT or BYTE data columns from a table. You want to store these columns together in the same cluster. |
| <i>cluster-key column list</i> | is one or more columns that provide a logical order for the TEXT or BYTE columns. The cluster-key columns must be from the same table as the TEXT or BYTE data columns. |

For example, suppose you are a sporting-goods distributor and you have photographs taken of your stock to include in an electronic catalog. With the Optical Subsystem, you create a table called **catalog**, which is similar to the **catalog** table provided with the **stores_demo** database but which stores the **cat_picture** column on an optical platter rather than a magnetic disk.

TEXT and BYTE Data Columns

In the CREATE TABLE statement for **catalog**, you define a BYTE column called **cat_picture** and use the IN portion of the column definition to place the column in the optical family **stock_photos**. For the complete syntax of the CREATE TABLE statement, see the *IBM Informix Guide to SQL: Syntax*.

The Cluster-Key Column

Next, you decide that you want to keep the photographs for a given manufacturer together to avoid unnecessary platter exchanges during retrieval. To keep photographs for a given manufacturer together, use the CREATE OPTICAL CLUSTER statement and make the **manu_code** column (also a column in the **catalog** table) the **cluster-key** column. Because **manu_code** contains nine unique values (ANZ, HRO, HSK, KAR, NKL, NRG, PRC, SHM, SMT), the photographs are stored in nine unique clusters, each containing the photographs for a single manufacturer. You assign a cluster size of 18,000 kilobytes to the cluster for the following reason: each photograph requires 60 kilobytes of storage, and you want to allow for as many as 300 photographs per manufacturer.

How Optical Clustering Occurs

Insert the first photograph into the database. As the insert operation begins, the Optical Subsystem checks whether the data is stored on an optical platter and whether an optical cluster exists for it. If so, the Optical Subsystem uses an internal optical-cluster table to find the current volume for the cluster-key value. If the cluster-key (**manu_code**) value for the first insert is HSK, the Optical Subsystem asks the optical-storage subsystem to reserve 18,000 kilobytes on the current volume for this cluster key. This first photograph is then written onto the volume.

The Optical Subsystem adds the cluster-key value, HSK, and the location of the HSK cluster to the cluster-size information that is already recorded in its internal optical-cluster table. The Optical Subsystem also records the fact that 60 kilobytes in the HSK cluster are now used. When the next unique cluster-key value, NRG for example, is inserted, the optical-storage subsystem reserves another 18,000 kilobytes on the current volume for the data associated with this cluster key. This process is repeated each time the optical-storage subsystem encounters a unique cluster-key value.

As HSK and NRG photographs fill their respective clusters, the internal optical-cluster table tracks the amount of space used in each cluster. Before a photograph is written to an optical platter, the Optical Subsystem compares the size of the object with the amount of space that remains in its assigned cluster. If the cluster is too full to receive the photograph, the optical-storage subsystem creates a new cluster on the volume.

If all space on the current volume is filled or reserved, the optical-storage subsystem finds the next available volume in the family (a stand-alone optical-storage subsystem prompts you to mount an empty platter), and creates a new cluster on it for the cluster-key value. When the new cluster is created, the internal optical-cluster table is updated with the new volume number and the new value for space that remains in the current cluster. If an additional volume is not available in the family, an error is returned.

Calculating Logical-Log Space for Optical Data

When you install the Optical Subsystem, you do not need to calculate any logical-log space for optical TEXT and BYTE data. Dynamic Server does not write optical TEXT and BYTE data to either the logical log or physical log. Optical TEXT and BYTE data is *not* written to the logical-log backup tapes when the logical logs are backed up.

For information on estimating the size of your logical log and physical log for tblspace TEXT and BYTE data, refer to your *Administrator's Guide*.

Support for Multiple Optical Virtual Processors

The Optical Subsystem supports multiple optical virtual processors (VPs). To execute optical I/O requests on multiple drives in parallel, configure multiple optical VPs. You can fine-tune the number of optical VPs based on your system configuration and usage.

You can set the number of optical VPs equal to or greater than the number of optical drives on a multiprocessor system. The optical VPs are not bound to a particular optical drive. An idle optical VP can service requests for I/O on any optical drive.

Determining the Number of Optical Virtual Processors

The following guidelines can help you determine the number of optical VPs to configure for your Optical Subsystem. The optimal number of optical VPs depends on the type of system usage.

Guideline One: The maximum number of optical VPs is equal to the number of clients performing optical I/O concurrently.

Guideline Two: However, if several clients are performing optical I/O concurrently, then use at least one optical VP per drive.

Scenario One: Clients are requesting writes on optical volumes in the same family. Multiple optical VPs for each drive allow additional clients to complete the initial setup while the first client is performing its I/O. The goal is to keep all the drives busy and maximize performance. If the optical objects are large, two optical VPs per drive are sufficient. If the optical objects are tiny (as in the case of signature verification), you need several optical VPs to keep the drives busy.

Scenario Two: Clients are requesting writes on the current optical volume in many different families. The write requests are randomly distributed across families. If the number of write requests is high, you would gain performance by having more optical VPs because you are more likely to get an I/O request for an already mounted platter before it is dismounted. Multiple optical VPs reduce the number of platter exchanges, a time-consuming operation.

Scenario Three: In a read-intensive environment, you should use one optical VP per optical drive.

Scenario Four: If the number of concurrent optical I/O request rates is low, a single optical VP for all optical drives might be adequate.

The Optical Subsystem API

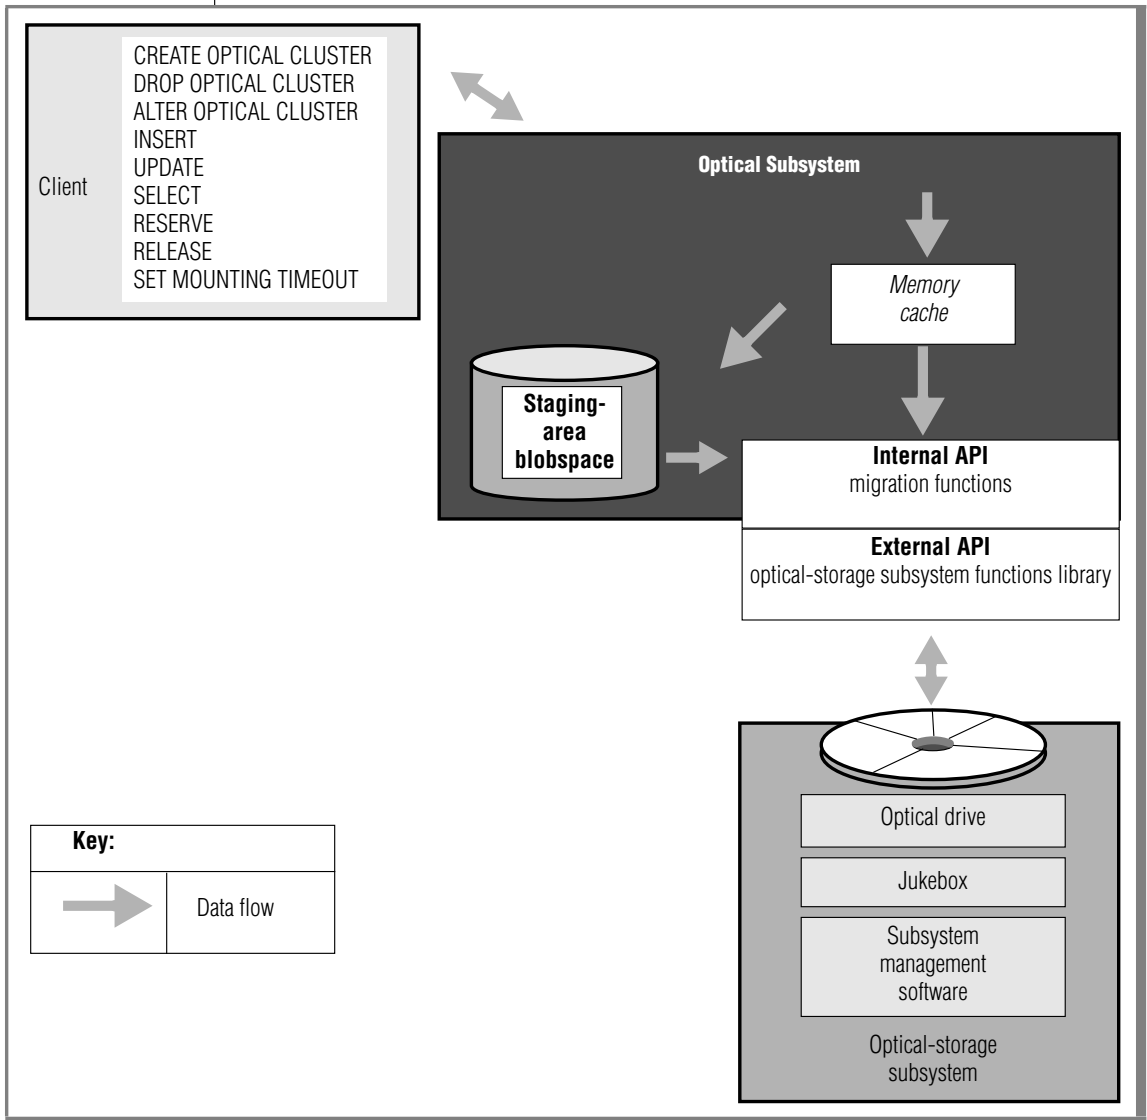
The application programming interface (API) between the Optical Subsystem and the optical-storage subsystem consists of the following layers:

- An *internal* layer that represents the interaction between the Optical Subsystem and the API
- An *external* layer that represents the interaction between the API and the optical-storage subsystem

The optical-storage subsystem vendor provides the external layer of the interface in the form of a shared library that is linked to the Optical Subsystem. For information on how to install the vendor library and how to link it with the Optical Subsystem, see [Chapter 2, “Installation and Initial Configuration.”](#)

Figure 1-3 illustrates the relationships between the Optical Subsystem, the API, and the optical-storage subsystem.

Figure 1-3
Optical Subsystem, the API, the Staging-Area Blobspace, the Memory Cache, and the Optical-Storage Subsystem



The Memory Cache and Staging-Area Blobspace

The staging-area blobspace improves the performance of writing TEXT and BYTE data that is larger than the cache size. The Optical Subsystem receives TEXT and BYTE data from the client application in 1-kilobyte pieces for a single row at a time. If the memory cache is full and cannot hold all pieces of TEXT and BYTE data, or if the cache is not being used, then the Optical Subsystem writes the TEXT and BYTE data to the staging-area blobspace. The Optical Subsystem uses the same memory cache for all client applications that are writing TEXT and BYTE data.

The configuration parameter **OPCACHEMAX**, located in the **ONCONFIG** file, specifies the size of the memory cache in kilobytes. If the value is 0, the Optical Subsystem does not use the cache. You can set the client environment variable, **INFORMIXOPCACHE**, to restrict the amount of memory cache that the client application uses. The value for **INFORMIXOPCACHE** is specified in kilobytes and can be any value less than or equal to **OPCACHEMAX**. If **INFORMIXOPCACHE** is not set, the client application can use as much of the memory cache as is available.

For more information about the **ONCONFIG** file, see the *IBM Informix Guide to SQL: Reference*. For more information about the **OPCACHEMAX** configuration parameter, see the *Administrator's Reference*.

The Optical Subsystem stores all TEXT or BYTE data for a given row in the memory cache until the last bit of TEXT or BYTE data in the row has been received or until the cache is full. If TEXT or BYTE data in the row does not fit in the memory cache, it is flushed out of the cache to the staging-area blobspace. If the next column of TEXT or BYTE data in the row fits in the available space in the memory cache, it stays in the memory cache.

For example, imagine that you have a memory cache of 150 kilobytes, that the setting for the configuration parameter `OPCACHEMAX` is 150 kilobytes, and that the environment variable `INFORMIXOPCACHE` is not set. Your current row has three columns of TEXT data of the following sizes: 30 kilobytes, 180 kilobytes, and 70 kilobytes. The first TEXT data column is 30 kilobytes and the Optical Subsystem writes it to the memory cache. The second TEXT data column is 180 kilobytes and the Optical Subsystem tries to use the memory cache but the data is too large. The Optical Subsystem writes the TEXT data from the second column to the staging-area blobSpace. The third column of TEXT data is 70 kilobytes and fits in the memory cache. The Optical Subsystem now outmigrates all three TEXT data columns in the row to the optical-storage subsystem in 50-kilobyte pieces in the order the TEXT data columns were processed: the first TEXT data from the memory cache; the second TEXT data from the staging-area blobSpace; and the third TEXT data from the memory cache.

If you want to outmigrate a different number of bytes than the default of 50 kilobytes, you can configure the outmigration amount using the storage manager for the optical-storage subsystem. [Figure 1-3 on page 1-13](#) illustrates the use of the memory cache and the staging area when TEXT and BYTE data is outmigrated to the optical-storage subsystem.

You create the staging-area blobSpace with the **onspaces** command line utility. Before you use the optical-storage subsystem, you must perform the following tasks:

- Edit the `ONCONFIG` file and set the `STAGEBLOB` parameter to the name of the staging-area blobSpace.
- Create the staging-area blobSpace.
- Restart the database server to enable the Optical Subsystem to recognize the optical-storage subsystem.

Before you use the optical-storage subsystem, you might also want to perform the following tasks:

- Adjust the configuration parameter `OPCACHEMAX` to something other than the default (128 kilobytes).
- Specify the size for the `INFORMIXOPCACHE` environment variable in client environments.

For detailed information for each step, see [“Creating the Staging Area” on page 2-5](#).

Memory Cache

Memory for the memory cache is allocated as it is used. If OPCACHEMAX is set at 500 (500 kilobytes) and the largest TEXT or BYTE data is 50 kilobytes, the Optical Subsystem uses only 50 kilobytes for the memory cache.

Because the setting for the configuration parameter OPCACHEMAX is system wide, 100 percent of the memory cache is available if you are the only user. If additional users compete for use of the memory cache, more data is written to the staging-area blobospace. To properly assess the use of the memory cache, look at how many applications are processing TEXT or BYTE data objects and set the OPCACHEMAX to accommodate the largest data object.

The system configuration parameter OPCACHEMAX uses a default value of 128 kilobytes; however, you can set OPCACHEMAX to any value. If OPCACHEMAX is set to 0, the memory cache is not used and the Optical Subsystem writes all TEXT and BYTE data to the staging-area blobospace.

The Staging Area

The structure of the staging-area blobospace is the same as all other blobspaces. When it is created, the staging area consists of only one chunk, but more can be added as desired. Blobospace free-map pages manage the space. For information on how to create a blobospace, see the *Administrator's Reference*.

The optimal size for the staging-area blobospace depends on the following application factors:

- The frequency of data storage
- The frequency of data retrieval
- The average size of the data

To calculate the size of the staging-area blobSpace, you must estimate the number of TEXT or BYTE data objects that you expect to reside in the staging-area blobSpace simultaneously and multiply that number by the average size. In a single transaction, all TEXT or BYTE data that moves to the staging area is held in the staging area until the transaction completes. The number of TEXT or BYTE data objects that reside in the staging area simultaneously depends on the number of rows that are inserted at one time. The minimum size of the staging-area blobSpace must be at least as large as the largest TEXT or BYTE data object that will reside in it. If the Optical Subsystem uses the staging-area blobSpace, it writes an entire TEXT or BYTE data object to the staging area before it outmigrates the object to the optical-storage subsystem.

If a hardware failure occurs in the staging-area blobSpace, the Optical Subsystem rolls the transaction back. The Optical Subsystem does not commit the transaction until it has written the data to the optical platter.

The staging-area blobSpace cannot be mirrored.

The STAGEBLOB Parameter

You specify the stageblob blobSpace with the STAGEBLOB configuration parameter. The STAGEBLOB parameter has two values separated by a comma. The first value gives the name of the blobSpace and the second value gives the number of optical VPs configured. In the following example, the *blobSpace_name* is the name of the stageblob blobSpace and 2 is the number of optical VPs configured. For information on configuring optical VPs, see [“Support for Multiple Optical Virtual Processors” on page 1-11](#).

```
STAGEBLOB blobSpace_name, 2
```

For information on how to set the STAGEBLOB parameter, refer to the *Administrator's Reference*.

The presence of the STAGEBLOB parameter in the ONCONFIG file signals the database server that an optical-storage subsystem is present. If the database server does not find the STAGEBLOB parameter in the ONCONFIG file or if STAGEBLOB has a blank tag, it behaves as if no optical-storage subsystem existed. If the STAGEBLOB parameter is present but you have not created the staging-area blobSpace, the database server displays the following message:

```
Invalid or missing name for Subsystem Staging BlobSpace
```

You must create the staging-area blobspace and reinitialize the database server before it can migrate TEXT or BYTE data to the optical-storage subsystem. For the procedure to create the staging-area blobspace, refer to [Chapter 2](#).

The onstat -O Option

You can use the **-O** option of the **onstat** utility to monitor available and allocated resources for the memory cache and the staging-area blobspace. [Figure 1-4](#) shows a sample of **onstat -O** output.

Figure 1-4
Example of the onstat -O Option

```
IBM Informix Dynamic Server, Version 9.4 UN105 --Online-- Up 00:45:18 -- 11656 Kbytes

Optical StageBlob Cache
System Cache Totals:
Size      Alloc. Avail.      Number   Kbytes
500       500     0           1       20
System Blob Totals:
Number   Kbytes
3        1500

User Cache Totals:
SID      User   Size      Number   Kbytes
94       doug   250        1       20
95       beth   500        0        0
User Blob Totals:
Number   Kbytes
1        300
2       1200
```

The example shows that Doug wrote one 20-kilobyte TEXT or BYTE data to the memory cache. He also tried to write a 300-kilobyte TEXT or BYTE data object, but it did not fit because he is limited to 250 kilobytes of the memory cache. Beth can use the entire 500-kilobyte memory cache, but the two TEXT or BYTE data objects that she tried to write did not fit in the memory cache (1,200 kilobytes) and were written to the staging-area blobspace. No optical writes are in progress in [Figure 1-4](#) (Avail = 0). No optical work was completed except for these two programs because the totals equal the sum of the two user entries (1,500 kilobytes). You can tell that the two programs are still connected because the entries are visible.

The API Internal Layer

The internal layer of the API manages the following operations:

- An *outmigration* operation that stores TEXT and BYTE data in the optical-storage subsystem
- An *immigration* operation that retrieves TEXT and BYTE data from the optical-storage subsystem

Outmigration is the operation through which the Optical Subsystem directs the optical-storage subsystem to store TEXT and BYTE data on an optical platter. Outmigration activity occurs in the following sequence:

1. The Optical Subsystem stores the TEXT or BYTE data in the memory cache in 1-kilobyte pieces.
2. The Optical Subsystem stores TEXT or BYTE data that exceeds the memory cache in the staging-area blobspace.
3. The Optical Subsystem makes the size and cluster information about the TEXT or BYTE data available to the optical-storage subsystem.
4. The optical-storage subsystem finds and reserves space to store the TEXT or BYTE data.
5. The Optical Subsystem transfers the TEXT or BYTE data to the optical-storage subsystem.
6. The Optical Subsystem signals the end of data migration. Once the outmigration process ends, the staging-area blobpages that held the migrating data are marked *free*. In addition, the memory cache is flushed and available for use.
7. The family number, volume number, and byte-offset that the optical-storage subsystem uses to describe the storage location are passed back to the Optical Subsystem, where the information is stored in the TEXT or BYTE data descriptor in the data row. (The TEXT or BYTE data descriptor resides in a database table column. The descriptor is the pointer to the actual TEXT or BYTE data on the optical-storage subsystem.)

Because data cannot be rewritten to the same location on a WORM optical platter, when the Optical Subsystem updates TEXT or BYTE data, it outmigrates it to a new, clean location on the platter and updates the descriptor to reflect the new location. The optical-storage subsystem cannot reclaim space on the WORM platter after any TEXT or BYTE data on the platter is modified or deleted.

Immigration is the operation through which the Optical Subsystem requests and gains access to TEXT or BYTE data that is stored in the optical-storage subsystem. Immigration activity occurs in the following sequence:

1. The Optical Subsystem requests access to TEXT or BYTE data by supplying the optical-storage subsystem with the family name, volume number, and byte offset.
2. The optical-storage subsystem locates the TEXT or BYTE data and retrieves it in pieces. It deposits each portion of the TEXT or BYTE data in memory buffers that the Optical Subsystem can access.
3. The optical-storage subsystem signals the end of optical data migration.

Memory for the Optical Transfer Buffer

The memory for the optical transfer buffer is allocated from a special pool called *opool*. The optical transfer buffer transfers BYTE and TEXT data to and from the Optical Subsystem buffer.

To track memory usage, issue the **onstat -g mem opool** command. [Figure 1-5](#) shows the output of the **onstat -g mem opool** command.

Figure 1-5
Example of the onstat -g mem opool Option

```
IBM Informix Dynamic Server Version 9.4.UN105--Online-Up 00:02:08--4968 Kbytes
Pool Summary:
  name  class  addr      totalsize  freesize  #allocfrag  #freefrag
  opool  V      a2ca018  8192      8072      1           1

Blkpool Summary:
  name  class  addr      size      #blks
```

The API External Layer

The external layer of the API consists of a set of low-level functions that implement the interaction between the API and the optical-storage subsystem. These functions are compiled in a library that the optical-storage subsystem vendor provides.

The Optical Subsystem Message File

Under certain exceptional conditions, the Optical Subsystem API writes messages to the message file. In general, the messages indicate that the Optical Subsystem made a request of the optical-storage subsystem and the optical-storage subsystem was unable to accomplish it. For the content of these messages and an appropriate action or response, if one is required, see the documentation from your optical-storage subsystem vendor. A message consists of two error numbers. The first error number is Informix specific, and the second is generated by the vendor's subsystem log file. The MSGPATH parameter in the ONCONFIG file specifies the UNIX pathname of the system message log file. For more information about the message log file, refer to the *Administrator's Reference*.

The database server writes the following messages concerning the Optical Subsystem to the message log file:

`oninit: invalid or missing name for Subsystem Staging Blobspace`

Cause: The STAGEBLOB configuration parameter specifies a blobspace, but the named blobspace does not exist. This message is normal the first time **oninit** is started with the Optical Subsystem. Check to see if you have accidentally created the STAGEBLOB as a dbspace instead of a blobspace.

Action: Use **onspaces** with the **-b** option to create the blobspace.

`Optical Subsystem is running`

Cause: You have set the value of the STAGEBLOB parameter in the configuration file, and the Optical Subsystem is communicating properly with the optical-storage subsystem.

Action: None required.

`Optical Subsystem is not running`

Cause: You have set the value of the STAGEBLOB parameter in the configuration file, but the Optical Subsystem cannot detect the existence of the optical-storage subsystem.

Or on UNIX, the `/usr/lib/iosm09a.so` link is not linked to the optical library.

Action: Check that the optical-storage subsystem is online.

`Optical Subsystem STARTUP Error`

Cause: The Optical Subsystem detects that the optical-storage subsystem is running, but the Optical Subsystem cannot communicate with it properly.

Action: Check your optical-storage subsystem for errors.

Unable to initiate communication with the Optical Subsystem

Cause: The optical driver that the optical-drive vendor supplies indicates that the drive is not accessible.

Action: Check the driver installation and the cabling between the computer and the drive.

Installation and Initial Configuration

In This Chapter	2-3
Prerequisites	2-4
Supporting the Shared Library	2-4
Creating the Staging Area	2-5
Naming the Staging Area on the STAGEBLOB Parameter	2-5
Configuring Optical Virtual Processors	2-6
Initializing the Optical Subsystem	2-6
Creating the Staging-Area BlobSpace	2-7
Verifying the Presence of the Optical Subsystem	2-9
Creating Optical Families	2-9
Testing the Connection	2-10
Test One	2-10
Test Two	2-10

In This Chapter

This chapter contains instructions for installing and configuring the Optical Subsystem. It contains the following sections:

- [“Prerequisites”](#) lists the procedures that your optical-storage subsystem vendor must complete before you can install the Optical Subsystem.
- [“Supporting the Shared Library”](#) provides detailed instructions on how to connect to the support library.
- [“Creating the Staging Area”](#) tells you how to set the STAGEBLOB parameters in the ONCONFIG configuration file and how to create the staging-area blob space.
- [“Creating Optical Families”](#) directs you to your vendor documentation for the optical-storage subsystem for instructions on how to create optical families and assign volumes to them.
- [“Testing the Connection”](#) directs you to perform two tests that indicate whether the connection between the Optical Subsystem and the optical-storage subsystem is operating correctly.

Prerequisites

You must obtain an optical-storage subsystem from an authorized optical-storage subsystem vendor before you can use your IBM Informix Optical Subsystem. Contact your sales representative to obtain a list of authorized optical-storage subsystem vendors.

Your vendor provides you with the support library for the optical-storage subsystem that is linked with the Optical Subsystem. You can link the Optical Subsystem to the support library during the installation process, or anytime after. Your optical-storage subsystem vendor also provides you with optical-storage subsystem utilities that perform the following functions:

- Monitor the operation of the optical-storage subsystem
- Introduce and initialize new optical platters of the optical-storage subsystem
- Perform administrative and management tasks related to the optical-storage subsystem

Your optical-storage subsystem vendor should also perform the following tasks before you install the Optical Subsystem:

- Install the optical-storage subsystem hardware
- Configure the host-operating system to support the optical-storage subsystem
- Run diagnostics to ensure proper operation of the optical-storage subsystem on the host computer

Supporting the Shared Library

The Optical Subsystem is installed automatically when you install your database server. For instructions on how to install the database server, see your *Installation Guide*.

If you are installing your database server for the first time, follow the instructions on how to configure and initialize the database server in your *Administrator's Guide*.

Edit your ONCONFIG file to set the `OPTICAL_LIB_PATH` to the shared optical library file used. For example, if you use Plexus optical server manager installed on the **d** drive of a Windows computer, you would set the following value in the ONCONFIG file:

```
OPTICAL_LIB_PATH d:\OSM\XDPSM\lib\libop.dll
```

Continue with [“Creating the Staging Area”](#) before you access the optical-storage subsystem.

Creating the Staging Area

When TEXT or BYTE data moves from the Optical Subsystem environment to the optical-storage subsystem, the data is first directed to a memory cache or a blob space, which is a *staging area*. Before using the optical-storage subsystem, the database server administrator must perform the following tasks:

- Name the staging-area blob space
- Initialize the database server
- Create the staging-area blob space

For more information, refer to your *Administrator’s Reference* and *Administrator’s Guide*.

Naming the Staging Area on the STAGEBLOB Parameter

You must provide the name of the staging-area blob space on the STAGEBLOB parameter.

On UNIX, the blob space name is the first field of the value of the STAGEBLOB parameter in the ONCONFIG file:

```
STAGEBLOB blob_space_name, num_optical_vps
```

◆

Windows

On Windows, the name of the staging-area blob space is the only value of the STAGEBLOB parameter in the ONCONFIG file:

```
STAGEBLOB blob space_name
```



UNIX

Configuring Optical Virtual Processors

On UNIX, you can configure more than one optical VP. To set the number of optical VPs, use the second field of the value of the STAGEBLOB parameter. For example, if you want to use three optical VPs, specify the following values:

```
STAGEBLOB myblob,3
```

Do not leave spaces between the comma following the name of the blob space and the number of optical VPs.

Initializing the Optical Subsystem

Initialize the database server shared memory. When you initialize the database server, the following error message appears on the screen:

```
Invalid or missing name for Subsystem Staging Blob space
```

The Optical Subsystem is still initialized even though the staging-area blob space has not yet been created. The staging-area blob space name supplied as the STAGEBLOB parameter informs the database server that an optical-storage subsystem is present. You must create the staging-area blob space, however, before you can outmigrate TEXT and BYTE data to the optical-storage subsystem.

UNIX

Creating the Staging-Area BlobSpace

The staging-area blobSpace is the same as any other blobSpace. The name of the staging-area blobSpace must be the same as the name that you provided on the STAGEBLOB parameter. After you create the staging-area blobSpace, you must restart the database server.

On UNIX, you create the staging-area blobSpace by using the **onspaces** utility.

To initialize the Optical Subsystem and create the staging-area blobSpace on UNIX

1. Name the staging-area blobSpace by editing the ONCONFIG file.
2. Log in as user **informix**.
3. Initialize the database server:
 - Go to the Control Panel and double-click the **Services** icon.
 - Select your database server name in the Services window and click **Start**.
 - If you are initializing the database server for the first time, type **-iy** in the Startup Parameters window. (Use the **-iy** parameters ONLY on new systems. The **-i** parameter deletes all existing databases).
4. Wait for the database server to have a status of online before you use **onspaces**. To check the status, use **onstat**.
5. Add staging-area blobSpace.
 - Change to the directory where the blobSpace should be stored:


```
cd pathname
```
 - Create the file for the blobSpace:


```
touch stageblob
```
 - Make sure the blobSpace has permissions 660 for user **informix**:


```
chmod 660 stageblob
onspaces -c -b stageblob -g 4 -p pathname/stageblob -o
0 -s 10000
```
6. Start new log:


```
onmode -l
```
7. Back up logs:


```
ontape -a
```

Windows

8. Do a level-0 archive:
`ontape -s -L 0`
9. Stop the database server:
`onmode -yuk`
10. Start the database server:
`oninit`



On Windows, you can use either the **onspaces** command-line utility to create the staging-area blobSpace.

To use onspaces to initialize the Optical Subsystem and create the staging-area blobSpace on Windows

1. Name the staging-area blobSpace by editing the ONCONFIG file.
2. Log in as user **informix**.
3. Initialize the database server:
 - Go to the Control Panel and double-click the **Services** icon.
 - Select your database server name in the Services window and click **Start**.
 - If you are initializing the database server for the first time, type **-iy** in the Startup Parameters window. (Use the **-iy** parameters ONLY on new systems. The **-i** parameter deletes all existing databases).
4. Wait for the database server to have a status of Started before you use **onspaces**. To check the status, use **onstat**.
5. Create a file for the blobSpace:
`copy NUL c:\ifmxdata\dbservername\stageblob`
6. Create the blobSpace:
`onspaces -c -b stageblob -g 4 -p
c:\ifmxdata\dbservername\stageblob -o 0 -s 10000`
7. Start new log:
`onmode -l`
8. Back up logs:
`ontape -a`

9. Do a level-0 archive:

```
ontape -s -L 0
```

10. Bring down **oninit**:

```
onmode -yuk
```

11. Start the *dbservername* service:

- Go to the Control Panel and double-click the **Services** icon.
- Select your database server name in the Services window and click **Start**.



Verifying the Presence of the Optical Subsystem

After configuration, check the message log file to make sure that your database server recognizes the Optical Subsystem. The following message should appear:

```
Optical subsystem is running.
```

If this message does not appear, check to make sure that the setup of the Optical Subsystem is correct and that the Optical Subsystem configuration is correct.

Creating Optical Families

TEXT and BYTE data that is outmigrated to the optical-storage subsystem is stored on a volume in an optical family that you assign using the CREATE TABLE statement. Before you can assign TEXT and BYTE data to an optical family, however, you must use the utilities that the vendor of your optical-storage subsystem provides to create an optical family and assign volumes to it. For the procedure to create an optical family, see the documentation that your optical-storage subsystem vendor provides.

Testing the Connection

The following tests ensure that you followed the relink and install procedures correctly and that the connection between the Optical Subsystem and the optical-storage system is operational.

Test One

Run the optical utilities of the optical-storage subsystem vendor that monitor the operation of the optical-storage subsystem and provide the options for platter management.

The Optical Subsystem log file should indicate no abnormalities.

Test Two

Start the database server while the optical-storage subsystem is up. Both the database server message file and the log file for the optical-storage subsystem should indicate normal operations.

The message `Optical Subsystem is running` appears in the message log if the optical-storage subsystem is operating normally. For the name of the log file for the optical-storage subsystem and the messages that should appear in it, see the vendor documentation for your optical-storage subsystem.

Using the Optical Subsystem

In This Chapter	3-3
Assigning TEXT and BYTE Columns to an Optical Platter	3-4
Reading TEXT and BYTE Data Columns from an Optical Platter	3-5
Clustering TEXT and BYTE Data.	3-5
Choosing the Cluster Key	3-6
Choosing the Cluster Size	3-7
Altering the Cluster Size	3-7
Dropping an Optical Cluster	3-8
Managing Volumes on the Optical Drives.	3-9
Reserving an Optical Volume	3-9
Releasing a Reserved Optical Volume.	3-10
Setting the Mounting Time-Out	3-10
Using TEXT and BYTE Data Descriptors	3-12
Locating the Optical Volume Where TEXT and BYTE Data Is Stored	3-14
Locating Columns Stored in Optical Families.	3-14
Migrating a Database with TEXT and BYTE Data on an Optical Platter	3-15
Using HPL to Unload and Load Optical Data	3-16
The dbexport and dbimport Utilities	3-16

In This Chapter

This chapter documents the SQL extensions that the Optical Subsystem supports to access WORM optical media. It provides examples that show you how to perform the following functions:

- Assign a TEXT or BYTE data column to an optical platter (CREATE TABLE).
- Create an optical cluster (CREATE OPTICAL CLUSTER).
- Reserve an optical volume (RESERVE).
- Release an optical volume (RELEASE).
- Set the mounting time-out (SET MOUNTING TIMEOUT).
- Locate the optical volume where a TEXT or BYTE column is stored (FAMILY() and VOLUME() functions).
- Use TEXT or BYTE data descriptors to enable multiple tables to share the same TEXT and BYTE data (DESCR() function).
- Migrate TEXT and BYTE data in an optical family from one Optical Subsystem to another (HPL, **dbexport** and **dbimport**).

The examples use the **cat_descr** and **cat_picture** TEXT and BYTE columns from the **catalog** table of the **stores_demo** database. You can use the optical SQL statements from DB-Access or from a program developed with one of the IBM Informix SQL APIs, such as IBM Informix ESQL/C.

Assigning TEXT and BYTE Columns to an Optical Platter

To tell the Optical Subsystem to store a TEXT or BYTE column on an optical platter, specify an optical *family name*, in the column-definition portion of the CREATE TABLE statement. The administrator for the optical-storage subsystem must create the optical family name in the optical-storage subsystem before you can execute a CREATE TABLE statement that refers to it. For the procedure to create the optical family name, see your subsystem vendor documentation.

In the following example, the TEXT column **cat_descr** of the **catalog** table is stored with the table and the BYTE column **cat_picture** is stored in the optical family **catalog_stores_demo**:

```
CREATE TABLE catalog
(
  catalog_num      SERIAL (10001),
  stock_num        SMALLINT NOT NULL,
  manu_code        CHAR(3) NOT NULL,
  cat_descr        TEXT IN TABLE,
  cat_picture      BYTE IN catalog_stores_demo,
  cat_advert       VARCHAR (255, 65)
)
IN dbspace10
```

When the statement is executed, the database server checks to determine whether **catalog_stores_demo** is a blob space or an optical family name. Either of the following conditions returns an error:

- The name refers to both a blob space *and* an optical family name.
- The name does not refer to either a blob space *or* an optical family name.

Once you create the table, you can insert TEXT and BYTE data in an optical family in the following ways:

- With the LOAD statement (DB-Access) or the **dbload** utility
- From locator variables (IBM Informix ESQL/C)

Reading TEXT and BYTE Data Columns from an Optical Platter

Use the SELECT statement to read a TEXT or BYTE data column from an optical platter in the same way that you read a TEXT or BYTE data column from a magnetic disk. The syntax of the SELECT statement does not change to support access to an optical-storage subsystem. However, the Optical Subsystem supports extensions to SQL that affect how efficiently a TEXT or BYTE data column is read from the optical platter.

- Clustering TEXT and BYTE data
- Managing optical volumes

Clustering TEXT and BYTE Data

You can use the CREATE OPTICAL CLUSTER statement to store logically related TEXT and BYTE data on the same volume of an optical platter. Storing related TEXT and BYTE data together minimizes platter exchanges when the TEXT and BYTE data is retrieved. The CREATE OPTICAL CLUSTER statement requires you to perform the following steps:

- Assign a name to the optical cluster.
- Specify the TEXT or BYTE data columns in a table that are to be stored together in the optical cluster.
- Specify the cluster key.

The cluster key consists of one or more columns in the table that define the logical grouping of the TEXT or BYTE data columns. The columns that make up the cluster key must not be TEXT or BYTE data columns.
- Specify the size of the cluster in kilobytes.

The following CREATE OPTICAL CLUSTER statement allocates a cluster of 18,000 kilobytes on an optical platter for the **cat_picture** BYTE column in the **catalog** table. It assigns the name **catalog_catpics** to the cluster.

```
CREATE OPTICAL CLUSTER catalog_catpics
FOR catalog (cat_picture)
ON (manu_code)
CLUSTERSIZE 18000
```

The cluster key for **catalog_catpics** is **manu_code**. Therefore the **cat_picture** TEXT and BYTE data that have the same value in the **manu_code** column are stored together on the same volume or volumes, if more than one is required.

Choosing the Cluster Key

In the preceding example, the **manu_code** column was chosen for the cluster key because, of all the columns in the **catalog** table, it best meets the following criteria for choosing a cluster key:

- The **manu_code** column provides a logical order for retrieving the TEXT and BYTE data.
- The **manu_code** column contains duplicate values.

The first criterion for a cluster key, which can consist of more than one column, is that it provides a logical order for retrieving the TEXT and BYTE data. The purpose of the cluster key is to optimize retrieval of the TEXT and BYTE data by grouping related TEXT and BYTE data on the same volume. Grouping TEXT and BYTE data by a cluster key is only advantageous, however, if the TEXT and BYTE data is normally retrieved in the order that it is grouped. If the **cat_picture** TEXT and BYTE data is grouped by **stock_num** and then retrieved by **manu_code**, as the following example shows, the optical-storage subsystem might still need to perform multiple platter exchanges to retrieve the TEXT and BYTE data:

```
SELECT cat_picture FROM catalog WHERE manu_code = 'HRO'
```

The second criterion for a cluster key is that it must contain duplicate values. Of all the columns in the **catalog** table, only the first three, **catalog_num**, **stock_num**, and **manu_code**, provide logical sequences for accessing the table. Of these columns, **catalog_num** is not useful because it has a SERIAL data type, which produces a unique value for each row in the table. Clustering on this column produces a cluster for every row in the table where **cat_picture** is not null.

Assume, for example, that 26 rows in the **catalog** table have pictures in the **cat_picture** column, and that the **manu_code** column has seven unique values with occurrences as the following example shows:

manu_code	occurrences
ANZ	7
HRO	6
KAR	1
NRG	1
PRC	4
SHM	5
SMT	2

For these characteristics, **manu_code** produces seven unique clusters.

Choosing the Cluster Size

The cluster size is the size of the cluster expressed in kilobytes. It is based on the number of TEXT and BYTE data objects that you expect to store in each cluster and the average size of the TEXT and BYTE data. You can calculate it as follows:

$$\text{cluster size} = \text{estimated number of TEXT and BYTE data objects} * \text{average TEXT and BYTE data size}$$

If the average size of the **cat_picture** images is 60 kilobytes, and you want to store up to 300 TEXT and BYTE data objects per cluster, the cluster size is 18,000 kilobytes.

Altering the Cluster Size

You can use the ALTER OPTICAL CLUSTER statement to change the size of an optical cluster for all clusters that are created after the statement is executed. In the following example, the size of the **catalog_catpics** cluster is reduced to 6,000 kilobytes, a cluster size that stores up to one hundred 60-kilobyte TEXT or BYTE data objects per cluster:

```
ALTER OPTICAL CLUSTER catalog_catpics CLUSTER SIZE 6000
```

The ALTER OPTICAL CLUSTER statement changes the cluster size only for TEXT or BYTE data objects created after the statement is executed. Clusters that were created previously remain their original size.

For more information, see [“ALTER OPTICAL CLUSTER” on page 4-4](#).

Dropping an Optical Cluster

You can use the DROP OPTICAL CLUSTER statement to terminate optical clustering. The following example terminates clustering for **catalog_catpics**:

```
DROP OPTICAL CLUSTER catalog_catpics
```

The DROP OPTICAL CLUSTER statement does not affect the TEXT and BYTE data objects that were already stored in a cluster. It terminates clustering only for the affected columns in the future.

You can also use the DROP OPTICAL CLUSTER statement with the CREATE OPTICAL CLUSTER statement to change either the TEXT or BYTE data columns or the cluster-key columns for an optical cluster. For example, you can change the cluster key for **catalog_catpics** from **manu_code** to **stock_num** with the following statements:

```
DROP OPTICAL CLUSTER catalog_catpics;  
CREATE OPTICAL CLUSTER catalog_catpics  
  FOR catalog (cat_picture)  
  ON (stock_num)  
  CLUSTER SIZE 3000
```

These statements do not affect TEXT and BYTE data objects that were previously clustered on **manu_code** because you cannot alter data written to a WORM optical platter. These statements change the clustering only for future inserts or updates. If you want to recluster the TEXT and BYTE data objects that were clustered by **manu_code**, you must update them. The TEXT and BYTE data objects are then rewritten to the new cluster. The TEXT or BYTE data descriptors are updated with the new location of the TEXT and BYTE data objects. The space that these TEXT and BYTE data objects originally occupied is lost because you cannot reuse space on WORM optical media.

For more information, see [“DROP OPTICAL CLUSTER” on page 4-10](#).

Managing Volumes on the Optical Drives

Your application can perform the following operations to manage the volumes on the optical drives and, consequently, affect the performance of the application:

- Reserve an optical volume to keep it mounted during a set of operations
- Release a reserve request to free an optical drive
- Set the time-out period for mounting an optical volume

Reserving an Optical Volume

You can use the RESERVE statement to keep a required optical volume mounted on an optical drive while you periodically access it during a set of operations. The RESERVE statement implies a request to mount the specified volume if it is not currently mounted.

When multiple users access the optical platters simultaneously, you might need to reserve a volume if your application accesses it consecutively. Otherwise, individual accesses are interleaved with requests from other applications so that each access by your application could require the volume to be remounted on the optical drive. Not reserving a volume can result in an unsatisfactory response time from the optical-storage subsystem.

For each RESERVE request, the optical-storage subsystem adds the value 1 to a reserve counter for that volume. If the value of the reserve counter is greater than zero, you cannot remove the volume. You can issue multiple RESERVE requests for a volume, but you must release each one before you can remove the volume from the drive. When you finish accessing the volume, use the RELEASE statement to release it. For details, see [“Releasing a Reserved Optical Volume.”](#)

If all the optical volumes mounted on the optical drives are reserved when you issue a RESERVE statement, the optical-storage subsystem waits for the number of seconds that the SET MOUNTING TIMEOUT statement specifies or for the default time-out period. For more information, see the [RESERVE](#) and [SET MOUNTING TIMEOUT](#) statements in [Chapter 4, “SQL for the Optical Subsystem.”](#)

You can use a SELECT statement to obtain the family name and volume number of the volume that you want to reserve. The FAMILY() and VOLUME() function expressions return the family name and volume number, respectively. In the following example, the SELECT statement uses the FAMILY() and VOLUME() functions to identify the volume that contains pictures of stock for manufacturer HRO. The RESERVE statement then reserves the volume, using the family name and volume number that the SELECT statement returned.

```
SELECT FAMILY(cat_picture), VOLUME(cat_picture)
  FROM catalog
  WHERE manu_code = 'HRO'
...
RESERVE 'catalog_1991' 013
```

Releasing a Reserved Optical Volume

You must release each RESERVE request for an optical volume with the RELEASE statement. The RELEASE statement cancels the latest RESERVE request for an optical volume by subtracting the value one (1) from the reserve counter for that volume. When the reserve counter is reduced to zero, you can remove the optical volume from the drive. The following statement releases volume 013 of the optical family:

```
RELEASE 'catalog_1991' 013
```

An error is returned if you issue a RELEASE statement for a volume that is not reserved.

The vendor of your optical-storage subsystem supplies a utility that you can also use to release a volume. The optical-storage subsystem administrator must release a volume with this utility when a program terminates prior to releasing a reserved volume.

Setting the Mounting Time-Out

The SET MOUNTING TIMEOUT statement controls the amount of time that your application waits for a volume to be mounted before it abandons the request. When you request access to an optical volume, the optical-storage subsystem cannot mount it immediately if all the optical drives are busy processing requests from other users. In addition, one or more drives might be reserved at the time of your request. If a SET MOUNTING TIMEOUT statement is not executed, the optical-storage subsystem times out after five minutes (300 seconds) if it cannot mount the requested volume.

The SET MOUNTING TIMEOUT statement provides the following options for setting the time-out period:

- Wait indefinitely until a drive is available for the requested volume
- Return the request for a volume immediately if it is not currently mounted on an optical drive
- Specify a number of seconds to wait for a volume to be mounted
- Wait for a volume to be mounted only if a drive is available

If it is mandatory that a particular optical volume be mounted before an application continues, it might be appropriate for the application to wait an indefinite amount of time for that volume to mount. If so, the following statement has that effect:

```
SET MOUNTING TIMEOUT TO WAIT
```

If multiple users and multiple applications are competing for the optical drives, you might not always want your application to wait for a volume that is not currently mounted. You might set the mounting time-out to NOT WAIT. For example, if your application retrieves data about an item of merchandise and the TEXT or BYTE data on an optical platter is desirable but not required, the application might not display the data if it is not immediately available.

The following example causes a request for a volume to return immediately if the volume is not currently mounted:

```
SET MOUNTING TIMEOUT TO NOT WAIT
```

Normally, you should set the time-out period to the number of seconds that is appropriate for the application to wait for a volume. The following statement sets the mounting time-out period to 30 seconds. If the volume is not mounted within 30 seconds, the process times out and control is returned to the application.

```
SET MOUNTING TIMEOUT TO 30
```

Set the mounting time-out period to zero seconds to cause the process to wait for a volume to be mounted only if a drive is currently available, as follows:

```
SET MOUNTING TIMEOUT TO 0
```

When you specify a time-out period of zero, the optical-storage subsystem waits for a volume to mount if a drive is currently available for it. When you specify the NOT WAIT option, the optical-storage subsystem does not wait for the requested volume to mount if it is not currently available on a drive.

For more information, see [“SET MOUNTING TIMEOUT” on page 4-16](#).

Using TEXT and BYTE Data Descriptors

Data rows that include TEXT or BYTE data do not include the TEXT or BYTE data in the row itself. Instead, the data row contains a 56-byte TEXT or BYTE data descriptor that includes a pointer to the location where the first segment of data is stored. The descriptor can point to a dbspace blobpage, a blobpage, or an optical-storage location.

You can use the DESCR() function expression to retrieve the TEXT or BYTE data descriptor for TEXT or BYTE data that is stored on WORM optical media. The DESCR() function returns the descriptor in an encoded format that is 112 bytes long. Using the DESCR() function as a column value in INSERT and UPDATE statements enables you to create additional pointers to existing TEXT and BYTE data objects on the optical media. Creating additional pointers to TEXT and BYTE data objects saves space on the optical platters by allowing multiple tables to refer to the same physical TEXT and BYTE data.

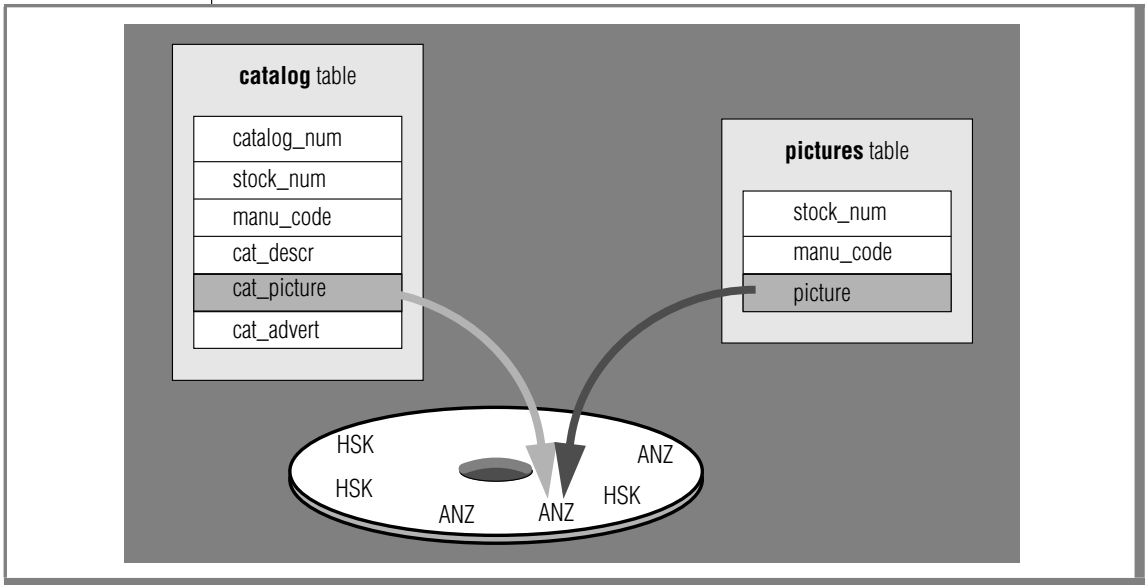
Before the TEXT or BYTE data descriptor is inserted in a data row, it is decoded and validated to verify consistency and legitimacy. The optical-storage subsystem performs all encoding, decoding, and validations through the Optical Subsystem calls to API functions. If the validation process fails, the insert operation fails.

The following example uses DESCR() in an INSERT statement to insert TEXT or BYTE data descriptors, or pointers to existing TEXT and BYTE data objects, in the following hypothetical **pictures** table. Assume that the **picture** column is a BYTE column assigned to an optical platter.

```
INSERT INTO pictures (stock_num, manu_code, picture)
  SELECT stock_num, manu_code, DESCR(cat_picture)
  FROM catalog
  WHERE manu_code = 'ANZ'
```

Figure 3-1 illustrates the result of the preceding INSERT statement for a row in the **pictures** table.

Figure 3-1
Two Descriptors That Point to the Same TEXT or BYTE Data Object on a WORM Platter



The following example uses DESCR() to update the **pictures** table with all pictures from the **cat_picture** column of the **catalog** table:

```
UPDATE pictures
SET (picture) =
  (SELECT DESCR(cat_picture)
   FROM catalog
   WHERE pictures.stock_num = catalog.stock_num
   AND pictures.manu_code = catalog.manu_code
   AND cat_picture IS NOT NULL)
```

You can use the DESCR() function expression for TEXT or BYTE data columns that are stored on an optical platter. An error is returned if you use DESCR() on a nonoptical TEXT or BYTE data column.

Locating the Optical Volume Where TEXT and BYTE Data Is Stored

You can use the FAMILY() and VOLUME() function expressions in a SELECT statement to obtain the family name and volume number where a TEXT and BYTE data object is stored. The SELECT statement in the following example produces a list of the optical volumes that hold pictures of bicycle helmets (stock_num = 110):

```
SELECT FAMILY(cat_picture), VOLUME(cat_picture)
FROM catalog
WHERE stock_num = 110
```

Locating Columns Stored in Optical Families

To determine which columns in a database are stored on an optical platter, query the **sysblobs** table. The **sysblobs** table contains four columns that provide the following information:

- The space name (blobspace, dbspace, or family name) where the column is stored
- The media type (M = magnetic; O = optical)
- The table identification number
- The column number within the table

The following SELECT statement shows the locations of all TEXT or BYTE data columns in a given database:

```
SELECT * FROM sysblobs
```

The output that the preceding SELECT statement produces for the **stores_demo** database indicates that column number 5 in the table with tabid 109 is stored in the optical family **family1**.

spacename	type	tabid	colno
rootdbs	M	109	4
family1	O	109	5

With this information, you can query the **syscolumns** table to obtain the name of the column, as follows:

```
SELECT colname FROM syscolumns WHERE tabid = 109 AND colno = 5
```

That particular query produces the following output:

```
colname  
cat_picture
```

The **oncheck** utility with the **-pD** option also tells you whether a TEXT or BYTE data column is stored on an optical platter. However, the information is displayed on a row-by-row basis. The following **oncheck** command displays information for rows in the **catalog** table of the **stores_demo** database:

```
oncheck -pD stores_demo:catalog
```

For more information about **oncheck**, refer to the *Administrator's Reference*.

Migrating a Database with TEXT and BYTE Data on an Optical Platter

When a database contains TEXT and BYTE data that is stored in an optical family, you can migrate the database from one Optical Subsystem to another using the following utilities:

- High-Performance Loader (HPL)
- **dbexport** and **dbimport**

If you are migrating a database that contains TEXT and BYTE data stored on an optical platter, check your vendor documentation for the compatibility requirements among optical-storage subsystems. For example, the destination optical-storage subsystem might need to include the same TEXT or BYTE data family name as the one that the source optical-storage subsystem uses.

The **ontape**, **onunload**, and **onload** utilities do not unload or load TEXT and BYTE data stored on optical media.

Using HPL to Unload and Load Optical Data

You can use the HPL utility to migrate optical TEXT and BYTE data. With HPL you can either load or unload a complete table or a partial table. For example, you can unload selected columns or rows that contain optical TEXT and BYTE data and insert them into a table in another database.

For instructions on how to use the HPL utility, see the *IBM Informix High-Performance Loader User's Guide*.

The dbexport and dbimport Utilities

The **dbexport** utility unloads a database into ASCII files; the **dbimport** utility creates a database from ASCII files. You can use the **-d** option of the **dbexport** utility to specify that only the TEXT or BYTE data descriptor is written for a TEXT or BYTE data column that is stored on an optical platter. Migrating the descriptor permits a TEXT and BYTE data object that is stored on an optical platter to be shared in multiple databases, eliminating the need to store duplicate copies of it on an optical platter. The TEXT or BYTE data descriptor contains information about the location and size of the data. The descriptor holds all the information that is necessary to retrieve a TEXT or BYTE data object from an optical platter. If you do not specify the **-d** option, **dbexport** exports both the descriptor and the TEXT or BYTE data object.

If you want a different cluster arrangement in the destination database and you can afford the space on an optical platter, you can choose to duplicate the TEXT and BYTE data in a different cluster.

The following **dbexport** command exports the **stores_demo** database to the **exstores_demo** directory, exporting only descriptors for tables that contain TEXT or BYTE data columns. (The **-c** instructs the program to ignore all errors except fatal errors, **-o** specifies the output directory, and **-q** suppresses the echo of SQL statements.)

```
dbexport -c -d -o exstores_demo -q stores_demo
```


TEXT or BYTE data descriptors remain valid from one Optical Subsystem to another because each volume on an optical platter contains an internal tracking label. In addition to TEXT or BYTE data, each WORM volume contains a label with the following information:

- Family name
- Family number (also referred to as the optical-storage subsystem serial number)
- Volume number
- Optical-storage-subsystem identifier

Each optical-storage-subsystem device, whether a stand-alone drive or jukebox, has a unique optical-storage-subsystem identifier. The optical-storage subsystem uses the identifier and the family name to uniquely identify the family within the optical-storage subsystem. Using the identifier and the family name becomes an important issue in a distributed database environment where more than one optical-storage subsystem might be available.

For instructions on how to use **dbexport** and **dbimport**, see the *IBM Informix Migration Guide*.

SQL for the Optical Subsystem

In This Chapter	4-3
ALTER OPTICAL CLUSTER	4-4
CREATE OPTICAL CLUSTER	4-6
DROP OPTICAL CLUSTER	4-10
RELEASE	4-12
RESERVE	4-14
SET MOUNTING TIMEOUT	4-16
Function Expressions	4-18
DESCR()	4-19
FAMILY()	4-19
VOLUME()	4-20

In This Chapter

This chapter describes the syntax and usage of the individual SQL statements that the Optical Subsystem provides to support optical-storage subsystems. The syntax of each statement is illustrated in a diagram that shows the sequence of the required and optional elements. For a description of the conventions used in these diagrams, see [“Syntax Conventions” on page 8](#) of the Introduction. The SQL statements presented in this chapter are used exclusively with optical-storage subsystems. For additional information on the context in which each statement is used, see [Chapter 3, “Using the Optical Subsystem.”](#)

This chapter describes the following SQL statements:

- ALTER OPTICAL CLUSTER
- CREATE OPTICAL CLUSTER
- DROP OPTICAL CLUSTER
- RELEASE
- RESERVE
- SET MOUNTING TIMEOUT

This chapter also describes the following function expressions:

- DESCR()
- FAMILY()
- VOLUME()

+


D/B

E/C

ALTER OPTICAL CLUSTER

Use the ALTER OPTICAL CLUSTER statement to alter the size of an optical cluster.

Syntax



Element	Purpose	Restrictions	Syntax
<i>cluster_name</i>	The name of the cluster to alter	The cluster name must already exist.	Identifier segment; see <i>IBM Informix Guide to SQL: Syntax</i>
<i>clustersize</i>	Size of the cluster, in kilobytes, to be reserved for each unique cluster-key value	Clustersize must be less than the volume size.	Literal Number segment; see <i>IBM Informix Guide to SQL: Syntax</i>
<i>owner</i>	The user name of the owner of the cluster	The specified name must be a valid user name.	Identifier segment; see <i>IBM Informix Guide to SQL: Syntax</i>

Usage

To alter an optical cluster, you must be the owner of the cluster, have the Index privilege on the table, or have the DBA privilege.

In the following example, the CREATE OPTICAL CLUSTER statement creates **cat_clstr**, an optical cluster of 6,000 kilobytes for the **cat_picture** column in the **catalog** table. For each unique value of **stock_num** in the table, **cat_clstr** stores the associated data for this column. The following example changes the amount of space allocated for each **cat_clstr** from 6,000 kilobytes to 8,000 kilobytes. This change affects all clusters created after the ALTER OPTICAL CLUSTER statement is executed. Clusters created prior to an ALTER OPTICAL CLUSTER statement maintain their original size.

```
CREATE OPTICAL CLUSTER cat_clstr
  FOR catalog (cat_picture)
  ON (stock_num)
  CLUSTERSIZE 6000
ALTER OPTICAL CLUSTER cat_clstr CLUSTERSIZE 8000
```

Related Information

Related statements: CREATE OPTICAL CLUSTER and DROP OPTICAL CLUSTER

+

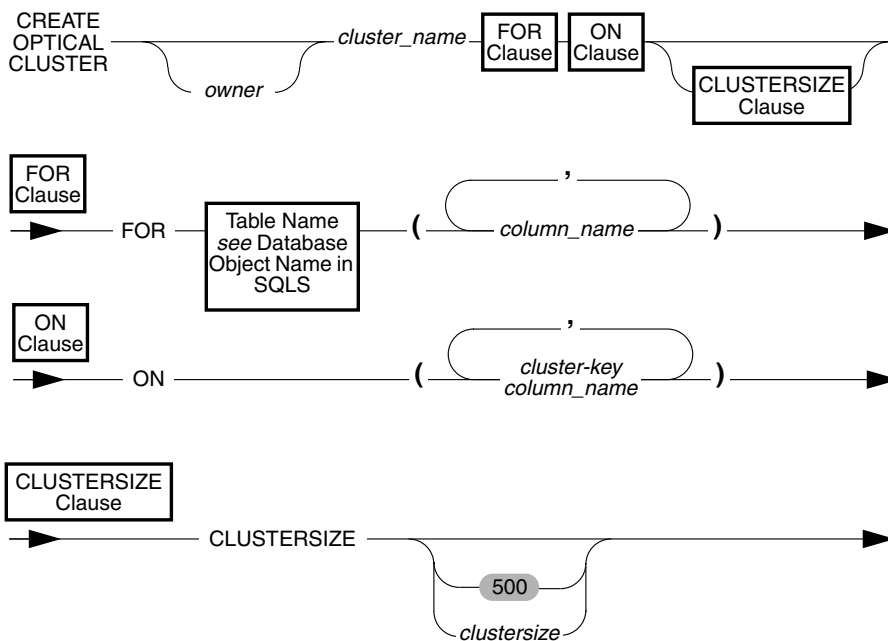
D/B

E/C

CREATE OPTICAL CLUSTER

Use the CREATE OPTICAL CLUSTER statement to create a cluster on an optical platter for one or more logically related TEXT or BYTE data columns.

Syntax



Element	Purpose	Restrictions	Syntax
<i>cluster-key</i> <i>column_name</i>	The name of the column that is the cluster key	You can specify up to 16 columns to create a composite cluster key. You cannot specify a TEXT or BYTE data column as part of a cluster key. The maximum length of a cluster key is 390 bytes.	Identifier segment; see <i>IBM Informix Guide to SQL: Syntax</i>
<i>cluster_name</i>	The name of the optical cluster	The name must be unique.	Identifier segment; see <i>IBM Informix Guide to SQL: Syntax</i>
<i>clustersize</i>	The amount of optical platter space to allocate for the cluster, specified in kilobytes If no clustersize is specified, the size of the cluster defaults to 500 kilobytes.	The clustersize must be less than the volume size.	Expression segment; see <i>IBM Informix Guide to SQL: Syntax</i>
<i>column_name</i>	The name of the TEXT or BYTE data column(s) to be clustered	You can assign up to 16 distinct columns to a cluster.	Identifier segment; see <i>IBM Informix Guide to SQL: Syntax</i>
<i>owner</i>	The valid user name of the cluster	The specified name must be a valid user name.	Identifier segment; see <i>IBM Informix Guide to SQL: Syntax</i>

Usage

An optical cluster enables logically related TEXT or BYTE data to be stored on the same volume, which minimizes time-consuming platter exchanges during retrieval.

You can create (and, by default, become owner of) an optical cluster if you are the owner of the table, have the Resource privilege on the database and the Index privilege on the table, or have the DBA privilege on the database.

Only a person with DBA privileges can create the optical cluster and specify another user as the owner of the cluster.

TEXT and BYTE Data Columns

The following restrictions are placed on the TEXT or BYTE data columns that are being clustered on the optical platter:

- All columns within a single clustering strategy must reside on the same optical family.
- A single column cannot be clustered more than once.

Cluster-Key Columns

A cluster key can be a single column (for example, **stock_num**) or a composite of up to 16 columns (for example, **stock_num**, **manu_code**). You cannot specify a TEXT or BYTE data column as a cluster key. The maximum length of a cluster key is 390 bytes. An error message is returned if you submit a cluster key that is longer than 390 bytes.

You can gain better performance if you organize the columns in a composite cluster key so that the column with the fewest number of duplicate values is placed first.

You can create more than one optical cluster for the same list of cluster-key columns. This option differs from the restrictions placed on indexes, each of which you must create from a unique set of index-key columns.

Clustersize

The *clustersize* is the size, in kilobytes, of the space reserved on an optical volume for each cluster with a unique cluster-key value. The size of the cluster can be enlarged beyond this size, if necessary, but it cannot exceed the size of a volume. If you do not specify *clustersize*, the default size of the cluster is 500 kilobytes.

Example

Assume that the CREATE TABLE statement for the **catalog** table in the **stores_demo** database specifies that the columns **cat_descr** and **cat_picture** are to be stored in the same optical family. In the following example, the CREATE OPTICAL CLUSTER statement creates **cat_clstr**, an optical cluster of 6,000 kilobytes. A new **cat_clstr** is allocated to store the data for **cat_descr** and **cat_picture** each time a row is inserted with a unique value for **manu_code**.

```
CREATE OPTICAL CLUSTER cat_clstr
  FOR catalog (cat_picture, cat_descr)
  ON (manu_code)
  CLUSTER SIZE 6000
```

If the **stores_demo** database contains six unique values for **manu_code**, the statement lays the groundwork for the optical-storage subsystem to create six optical clusters, one for each value. If the table contains ten rows for one **manu_code**, then the **cat_clstr** for that value contains the **cat_descr** and **cat_picture** data for all ten rows.

To create the cluster, the optical-storage subsystem reserves 6,000 kilobytes on the current volume. If the current volume is too full to accept the space reservation, the cluster is created on the next available volume in the family. The location and size of each cluster on the **manu_code** cluster key is tracked in an internal optical-cluster table. The name of the optical cluster, the number of its TEXT and BYTE data columns, and the cluster-key column numbers are stored in the **sysopclstr** table.

Related Information

Related statements: ALTER OPTICAL CLUSTER and DROP OPTICAL CLUSTER

For more information, see the CREATE TABLE statement in *IBM Informix Guide to SQL: Syntax*.

+


D/B

E/C

DROP OPTICAL CLUSTER

Use the DROP OPTICAL CLUSTER statement to end optical clustering for the TEXT or BYTE data columns that are associated with the cluster name.

Syntax



Element	Purpose	Restrictions	Syntax
<i>cluster_name</i>	The name of the optical cluster	The cluster name must already exist.	Identifier segment; see <i>IBM Informix Guide to SQL: Syntax</i>
<i>owner</i>	The user name of the owner of the cluster	The specified name must be a valid user name.	Identifier segment; see <i>IBM Informix Guide to SQL: Syntax</i>

Usage

The DROP OPTICAL CLUSTER statement terminates clustering on the optical platter for the TEXT or BYTE data columns associated with *cluster name* and drops the internal optical cluster table. After the DROP OPTICAL CLUSTER statement is executed, the data items that are associated with the cluster name are stored on the current volume; they are stored in the sequence that they are output to the optical-storage subsystem.

You must be the owner of the optical cluster or have the DBA privilege on the database to drop the optical cluster.

To change the cluster-key columns of an optical cluster, drop the old optical cluster and create a new optical cluster with the new cluster-key columns. You can change your optical clustering strategy so that the **cat_picture** and **cat_descr** columns of the **catalog** table are clustered according to **stock_num** instead of **manu_code**, as follows:

```
DROP OPTICAL CLUSTER cat_clstr
CREATE OPTICAL CLUSTER cat_stock_clstr
  FOR catalog (cat_picture, cat_descr)
  ON (stock_num)
  CLUSTERSIZE 6000
```

After the new optical cluster, **cat_stock_clstr**, is created on **stock_num**, all data that is inserted into the database for the **cat_picture** and **cat_descr** columns is optically clustered according to the associated value of **stock_num**. However, the new cluster does not alter the clustering of TEXT and BYTE data already stored on the WORM optical platters. A new cluster affects only subsequent outmigrations. Therefore, in this example, the TEXT and BYTE data that **cat_clstr** previously clustered retains its original locations. In subsequent outmigrations, the data for the **cat_descr** and **cat_picture** columns is clustered by **cat_stock_clstr**, based on the value of **stock_num**. Changing the cluster key impacts how the data is stored on the optical platter but not the integrity of the data.

Changing the clustering arrangement often can have an adverse affect on how efficiently data is retrieved. The TEXT and BYTE data that remains clustered by **manu_code** is effectively unclustered when **stock_num** accesses it. Accessing it by **stock_num** can result in frequent platter exchanges. To alleviate this problem, update the TEXT and BYTE data that **manu_code** previously stored so that it is rewritten in the new **stock_num** clusters. Because the optical media is WORM, however, the space where the TEXT and BYTE data was stored originally cannot be reused.

Administrators and programmers should be aware that the optical clustering strategy significantly affects the use of space on the optical platters.

Related Information

Related statements: CREATE OPTICAL CLUSTER

+


D/B

E/C

RELEASE

Use the RELEASE statement to cancel a reserve request for an optical volume.

Syntax



Element	Purpose	Restrictions	Syntax
<i>family_name</i>	A quoted-string constant that specifies a family name in the optical-storage subsystem	Must be an existing family name.	Quoted String segment; see the <i>IBM Informix Guide to SQL: Syntax</i> .
<i>family-name variable</i>	A CHARACTER or VARCHAR host variable that contains a family name	Must be an existing variable name.	Variable name must conform to language-specific rules for variable names.
<i>procedure-name variable</i>	The name of a variable defined in an SPL routine	Must be defined in the procedure and valid in the statement block.	Identifier segment; see the <i>IBM Informix Guide to SQL: Syntax</i> .
<i>volume_number</i>	The volume being released; specified as an integer	Must be reserved and mounted correctly.	Number must be an integer.

Usage

Each time that a RELEASE statement is executed, the value 1 is subtracted from the reserve counter for the volume. When the value of the reserve counter is zero, you can unmount the volume.

The volume specified in the RELEASE statement must be reserved and the platter that contains the volume must be mounted currently. An error is returned if the platter is not mounted or if the value of the reserve counter for the volume is zero.

You can issue multiple reserve requests, but you must release each one before you can remove the volume from the optical drive.

The following example releases a reservation for volume 013 in the **catalog_1991** optical family:

```
RELEASE 'catalog_1991' 013
```

Related Information

Related statements: RESERVE

+


D/B

E/C

RESERVE

Use the RESERVE statement to keep a particular optical volume mounted on the optical drive for the duration of a set of operations. The RESERVE statement implies a request to mount the specified volume if it is not currently mounted.

Syntax



Element	Purpose	Restrictions	Syntax
<i>family name</i>	A quoted-string constant that specifies a family name in the optical-storage subsystem	Must be a unique name.	Quoted String segment; see the <i>IBM Informix Guide to SQL: Syntax</i> .
<i>family-name variable</i>	A CHARACTER or VARCHAR host variable that contains a family name	Must be a unique variable name.	Variable name must conform to language-specific rules for variable names.
<i>procedure-name variable</i>	The name of a variable defined in an SPL routine	Must be defined in the procedure and valid in the statement block.	Identifier segment; see the <i>IBM Informix Guide to SQL: Syntax</i> .
<i>volume_number</i>	The number of the volume being reserved; specified as an integer	None.	Number must be an integer.

Usage

If the platter that contains the specified volume is not currently mounted on a drive, the optical-storage subsystem waits for a drive to become available and mounts the platter before it continues. The default waiting period for a mount attempt is five minutes (300 seconds). You can change this default period with the SET MOUNTING TIMEOUT statement.

You can issue multiple reserve requests for the same volume. Each time a RESERVE statement is executed for a volume, the value one (1) is added to the reserve counter for that volume. If no reserve requests have been issued for an optical volume, the value of the reserve counter is zero, indicating that you can remove the volume from the optical drive.

When you finish accessing the reserved volume, you can release it with the RELEASE statement. Each time a RELEASE statement is executed, the value 1 is subtracted from the reserve counter for that volume. Thus, you can issue multiple reserve requests, but you must release each one before you can remove the volume from the optical drive.

You can use a SELECT statement with the FAMILY() and VOLUME() functions to obtain the family name and volume number of the volume that you want to reserve. Then, you can specify that optical volume in the RESERVE statement, as follows:

```
SELECT FAMILY(cat_picture), VOLUME(cat_picture)
      WHERE manu_code = HRO AND stock_num = 301
...
RESERVE 'catalog_1991' 013
```

Related Information

Related statements: RELEASE

+

D/B


E/C

SET MOUNTING TIMEOUT

When you request access to an optical volume, the optical-storage subsystem cannot mount it immediately if all the optical drives are busy processing requests from other users. In addition, one or more drives might be reserved at the time of your request.

Use the SET MOUNTING TIMEOUT statement to control the amount of time that your application waits for a volume to be mounted before it abandons the request to continue processing. The mounting time-out disregards the amount of time spent waiting for an optical VP to become available. If the SET MOUNTING TIMEOUT statement is not executed, the default value for a mounting time-out is five minutes (300 seconds).

Syntax



Element	Purpose	Restrictions	Syntax
<i>seconds</i>	The maximum number of seconds the process should wait for access to a volume No default value exists for time-out within the statement.	Must be specified as a positive integer or zero. Mounting time of zero seconds is the same as NOT WAIT.	Literal Number segment; see <i>IBM Informix Guide to SQL: Syntax</i> .

Usage

If you set the time-out value to WAIT, all requests for a volume wait until the volume is mounted on an optical drive. The process that made the request waits indefinitely.

If you set the time-out value to NOT WAIT, all requests for a volume return immediately if that volume is not currently mounted on an optical drive.

If you set the time-out value to a positive integer, the integer specifies the number of seconds that the process waits for a volume to be mounted before the request returns back.

If you set the time-out value to zero, the process that made the request will wait for a volume to be mounted only if an optical drive is currently available. If all the drives contain either a reserved volume or a volume that is being accessed, the request returns back. If a drive is available, the process waits for the requested volume to be mounted. In general, mounting a volume takes about 15 seconds.

The following example sets the mounting time-out period to 45 seconds. When the application subsequently requests that a volume be mounted, it waits for up to 45 seconds for that volume to be mounted. If that volume cannot be mounted within the 45-second period, the process times out.

```
SET MOUNTING TIMEOUT TO 45
```

+

D/B

E/C

Function Expressions


The following function expressions support the optical-storage subsystem. Use them with the SELECT statement.

- DESCR()
- FAMILY()
- VOLUME()

You can use a function expression to perform various operations on column data or to obtain information from the database server about the contents of one or more columns. A function expression requires an argument or parameter. For all three of these function expressions, the argument is the name of a column that is stored on an optical platter. For more information on function expressions and a complete list of other function expressions, see the *IBM Informix Guide to SQL: Syntax*.

Use the DESCR(), FAMILY(), and VOLUME() function expressions to obtain information about TEXT or BYTE data columns that are stored on optical platters.

Syntax



Element	Purpose	Restrictions	Syntax
<i>column name</i>	The name of the TEXT or BYTE data column on which the function operates	The column must be stored on an optical platter.	Identifier segment; see <i>IBM Informix Guide to SQL: Syntax</i> .

Usage

The DESCR() function returns TEXT data that is the encoded descriptor for the TEXT or BYTE data column. The FAMILY() function returns the name of the optical family in which a column is stored. The VOLUME() function returns the number of the volume on which a column is stored. All three functions are valid *only* for data stored on an optical platter. They are valid whether the TEXT and BYTE data is stored sequentially or in clusters. An error message is returned if the functions are used on a column that is not stored on an optical platter.

DESCR()

The DESCR() function returns TEXT data, which is the encoded descriptor for the TEXT or BYTE data column. The encoded descriptor is 112 bytes. Use it as a column value in INSERT and UPDATE statements as a way to create additional pointers to existing TEXT and BYTE data on the optical platter. Before the descriptor is inserted in a data row, it is decoded and validated to verify its consistency and legitimacy. The optical-storage subsystem performs all encoding, decoding, and validations through calls by the Optical Subsystem to API functions. If the validation fails, the insert operation fails.

The DESCR() function is designed for WORM optical media only.

In the following example, the DESCR() function is used in an INSERT statement to fill the **picture** column of the **pictures** table with pointers to existing TEXT and BYTE data from the **catalog** table:

```
INSERT INTO pictures (stock_num, picture)
  SELECT stock_num, DESCR(cat_picture) FROM catalog
 WHERE manu_code = 'HRO'
```

FAMILY()

The FAMILY() function returns the name of the optical family in which a TEXT or BYTE data column is stored.

In the following example, the FAMILY() function obtains the *family name* of the optical platters that contain the data for the **cat_picture** column:

```
SELECT FAMILY(cat_picture) FROM catalog WHERE manu_code = 'HRO'
```

One row of output (*family name*) is generated for each row where **manu_code** is equal to HRO.

VOLUME()

The **VOLUME()** function returns the number of the volume where a TEXT or BYTE data column is stored.

Volumes are filled numerically. Once a volume is filled, or no room is available for a cluster, all subsequent writes occur on volumes with a higher number. When TEXT and BYTE data is stored sequentially, the highest volume number contains the most recently stored objects. However, this generalization is not true for clustered TEXT or BYTE data.

In the following example, the **VOLUME()** function generates a list of the volumes that contain data for the **cat_picture** column:

```
SELECT VOLUME(cat_picture) FROM catalog
```

Related Information

For more information on how to use **DESCR()**, **FAMILY()**, and **VOLUME()**, see [Chapter 3](#).

For the following information, see the *IBM Informix Guide to SQL: Syntax*:

- A complete list of other function expressions
- The syntax of the column-definition portion of the CREATE TABLE statement, which is used to assign a TEXT or BYTE data column to an optical family
- The SELECT statement

Notices

IBM may not offer the products, services, or features discussed in this document in all countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. (enter the year or years). All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

AIX; DB2; DB2 Universal Database; Distributed Relational Database Architecture; NUMA-Q; OS/2, OS/390, and OS/400; IBM Informix®; C-ISAM®; Foundation.2000™; IBM Informix® 4GL; IBM Informix® DataBlade® Module; Client SDK™; Cloudscape™; Cloudsync™; IBM Informix® Connect; IBM Informix® Driver for JDBC; Dynamic Connect™; IBM Informix® Dynamic Scalable Architecture™ (DSA); IBM Informix® Dynamic Server™; IBM Informix® Enterprise Gateway Manager (Enterprise Gateway Manager); IBM Informix® Extended Parallel Server™; i.Financial Services™; J/Foundation™; MaxConnect™; Object Translator™; Red Brick Decision Server™; IBM Informix® SE; IBM Informix® SQL; InformiXML™; RedBack®; SystemBuilder™; U2™; UniData®; UniVerse®; wintegrate® are trademarks or registered trademarks of International Business Machines Corporation.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Windows, Windows NT, and Excel are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names used in this publication may be trademarks or service marks of others.

Index

A

ALTER OPTICAL CLUSTER
 statement 3-7, 4-4
 ANSI compliance
 -ansi flag Intro-10
 level Intro-16
 Application programming interface (API)
 external layer 1-12, 1-21
 internal layer 1-12, 1-19
 Assigning a TEXT or BYTE column
 to optical disk 1-7, 3-4
 Autochanger 1-5

B

BLOB data type 1-3
 Blobspace, creating the staging-
 area 2-7
 Boldface type Intro-6
 BYTE column
 in catalog table 1-9, 3-4
 storing on optical disk 3-4
 BYTE data type, multiple
 copies 3-16

C

CLOB data type 1-3
 Cluster key
 choosing the 3-6
 composed of 3-5
 CREATE OPTICAL CLUSTER
 statement 1-9, 4-8
 defining 1-9

 specifying 4-8
 Cluster size
 altering 3-7, 4-4
 calculating 3-7, 4-8
 specifying 3-5, 4-6
 Clustered storage 1-7
 Clustering TEXT and BYTE
 columns, how it occurs 1-10
 Clustering TEXT and BYTE
 objects 3-5
 Code set, ISO 8859-1 Intro-4
 Column-definition portion of
 CREATE TABLE statement 1-7,
 3-4
 Column, TEXT or BYTE
 assigning to optical family 1-7
 changing 3-8
 CREATE OPTICAL CLUSTER
 statement 1-9
 defining 1-9
 Comment icons Intro-7
 Compliance
 with industry standards Intro-16
 Configuration parameter,
 OPCACHEMAX 1-14
 Contact information Intro-16
 Conventions,
 documentation Intro-5
 CREATE OPTICAL CLUSTER
 statement
 cluster key 4-8
 clustering TEXT and BYTE
 columns 1-9
 clustering TEXT and BYTE
 objects 3-5
 clustersize 4-8
 example 3-5, 4-9

logically ordering TEXT and
 BYTE columns 1-9
 requirements 3-5
 restrictions on TEXT and BYTE
 columns 4-8
 syntax and use 4-6
 CREATE TABLE statement
 assigning a column to optical
 platter 1-7
 assigning a TEXT or BYTE column
 to optical disk 1-7, 3-4
 creating optical family name 3-4
 defining TEXT or BYTE column
 for catalog table 1-9
 example 3-4
 IN clause 3-4

D

data descriptor, multiple
 copies 3-16
 Data types
 BLOB 1-3
 CLOB 1-3
 TEXT and BYTE 1-3
 DB-Access utility Intro-4
 dbexport 3-16
 dbimport 3-16
 Default locale Intro-4
 Demonstration databases Intro-4
 DESCR() function expression
 in an INSERT statement 3-12
 syntax and use 4-18
 Descriptor, TEXT and BYTE data
 DESCR() function
 expression 3-12
 using 3-12
 Documentation conventions
 sample-code Intro-12
 typographical Intro-6
 Documentation notes Intro-14
 Documentation notes, program
 item Intro-15
 Documentation, types of Intro-13
 documentation notes Intro-14
 machine notes Intro-14
 release notes Intro-14

DROP OPTICAL CLUSTER
 statement 3-8, 4-10

E

Environment variables Intro-6
 Environment variable,
 INFORMIXOPCACHE 1-14
 en_us.8859-1 locale Intro-4
 ESQ/C and the Optical
 Subsystem 3-3
 Extension to SQL, symbol
 for Intro-10

F

FAMILY() function expression
 locating data 3-14, 4-19
 Family, description of 1-6
 Feature icons Intro-7
 Features of this product,
 new Intro-5
 finderr utility Intro-15
 Function expressions
 DESCR() 4-19
 FAMILY() 4-19
 VOLUME() 4-20

G

Global Language Support
 (GLS) Intro-4

H

Help Intro-13

I

Icons
 feature Intro-7
 Important Intro-7
 platform Intro-7
 product Intro-7
 syntax diagram Intro-9
 Tip Intro-7

Warning Intro-7
 Important paragraphs, icon
 for Intro-7
 Industry standards, compliance
 with Intro-16
 INFORMIXDIR/bin
 directory Intro-5
 INFORMIXOPCACHE
 environment variable 1-14
 Immigration
 activity 1-20
 definition of 1-19
 Inserting TEXT and BYTE data
 descriptors 3-12
 Installation
 creating optical families 2-9
 creating the staging-area
 blobspace 2-5
 STAGEBLOB parameter 2-5
 testing the connection 2-10
 verifying that database server
 recognizes optical
 subsystem 2-9
 ISO 8859-1 code set Intro-4

J

Jukebox 1-5

L

Locale Intro-4
 default Intro-4
 en_us.8859-1 Intro-4
 Long identifiers Intro-5

M

Machine notes Intro-14
 Management of optical-storage
 subsystem 1-5
 Memory cache
 how space is allocated 1-16
 how to assess size 1-16
 monitoring available and
 allocated resources 1-18
 purpose 1-14

relationship to staging-area blobspace 1-14
 setting for client with INFORMIXOPCACHE 1-14
 setting for system with OPCACHEMAX 1-14
 Message file for error messages Intro-15
 Migrating a database, with TEXT and BYTE data on optical platter 3-15
 Multiple optical VPs
 guidelines for determining number of 1-11
 support for 1-11

N

New features of this product Intro-5

O

Online help Intro-13
 Online manuals Intro-13
 Onstat utility, using to monitor resources 1-18
 OPCACHEMAX configuration
 parameter 1-14
 Optical drive, performance 3-9
 Optical family 3-4
 Optical media
 accessing 1-5
 advantages 1-4
 platter 1-5
 storing TEXT and BYTE data 1-4
 Optical platter
 component of subsystem 1-5
 exchanges 3-5
 Optical Subsystem
 description of Intro-3, 1-3
 prerequisites for installation 2-4
 virtual processors 1-11
 Optical-storage subsystem
 assigning a TEXT or BYTE column to an optical family 1-7
 components 1-5

interaction between the API and the subsystem 1-12
 management software 1-5
 relationship with Optical Subsystem and the API 1-13
 stand-alone 1-5
 vendor library 1-12
 OPTICAL_LIB_PATH 2-5
 Outmigration
 activity 1-19
 definition of 1-19

P

Performance, application 3-9
 Platform icons Intro-7
 Product icons Intro-7
 Program group
 Documentation notes Intro-15
 Release notes Intro-15

R

Release notes Intro-14
 Release notes, program item Intro-15
 RELEASE statement
 example 3-10
 reserve counter 3-10
 syntax and use 4-12
 RESERVE statement
 FAMILY() function expression 3-10
 reserve counter 3-9
 syntax and use 4-14
 VOLUME() function expression 3-10

S

Sample-code conventions Intro-12
 SELECT statement
 obtaining family name 3-10
 obtaining volume number 3-10
 reading a TEXT or BYTE column 3-5

with FAMILY() function expression 3-14
 with VOLUME() function expression 3-14
 Sequential storage, TEXT and BYTE data 1-7
 SET MOUNTING TIMEOUT
 statement
 number of seconds option 3-11
 syntax and use 4-16
 using 3-10
 Shelf 1-5
 Software dependencies Intro-4
 SQL extensions, purpose 3-5
 SQL statements, using Optical Subsystem 3-3
 STAGEBLOB parameter
 blobspace_name 1-17
 naming the staging-area blobspace 2-5
 Staging-area blobspace
 creating 2-7
 creating with ON-Monitor or onspaces 1-15
 how to assess size 1-16
 using onstat utility to monitor resources 1-18
 stores_demo database Intro-4
 superstores Intro-4
 superstores_demo database Intro-4
 Syntax conventions
 elements of Intro-9
 example diagram Intro-11
 icons used in Intro-9
 sysblobs system catalog table 3-14
 sysopclstr table 4-9
 System catalog table, sysblobs 3-14

T

TEXT and BYTE data
 clustered storage 1-7
 data descriptor 3-16
 description of 1-3, 1-4, 1-11
 sequential storage 1-7
 storage location 1-19

TEXT column
 in catalog table 3-4
 storing on optical disk 3-4
 Tip icons Intro-7

U

UNIX operating system
 default locale for Intro-4
 Users, types of Intro-3
 Utility programs
 dbexport 3-16
 dbimport 3-16
 oncheck 3-15
 onload 3-15
 onspaces 1-15
 ontape 3-15
 onunload 3-15

V

Virtual processors (VPs) 1-11
 Volume
 description of 1-6
 managing on the optical drive 3-9
 SET MOUNTING TIMEOUT
 statement 3-10
 VOLUME() function expression
 identifying an optical
 volume 3-10
 syntax and use 4-20

W

Warning icons Intro-7
 Windows
 default locale for Intro-4
 Write-once-read-many
 (WORM) 1-4

X

X/Open compliance level Intro-16