

InterBase 6

API Guide



Borland/INPRISE

100 Enterprise Way, Scotts Valley, CA 95066 <http://www.interbase.com>

Inprise/Borland may have patents and/or pending patent applications covering subject matter in this document. The furnishing of this document does not convey any license to these patents.

Copyright 1999 Inprise/Borland. All rights reserved. All InterBase products are trademarks or registered trademarks of Inprise/Borland. All Borland products are trademarks or registered trademarks of Inprise/Borland. Other brand and product names are trademarks or registered trademarks of their respective holders.

1INT0055WW21001 6E1R0699

Table of Contents

<i>List of Tables</i>	<i>xiii</i>
---------------------------------	-------------

CHAPTER 1	Using the API Guide	
	Who should use this guide	17
	Topics covered in this guide	18
	Sample database and applications	19

Part I: API User's Guide

CHAPTER 2	Application Requirements	
	Requirements for all applications	24
	Including <code>ibase.h</code>	24
	Database requirements	24
	Transaction requirements	25
	Additional requirements	26
	Microsoft Windows requirements	26
	DSQL requirements	27
	Blob requirements	27
	Array requirements	28
	Event requirements	29
	Error-handling requirements	29
	Services requirements	30
	Compiling and linking	30

CHAPTER 3	Programming with the InterBase API	
	Basic procedure for application development	31
	Supported development environments	32
	User name and password requirements	32
	Specifying user name and password	33
	Using environment variables	33
	Setting a default database directory	34
	Setting a user name and password	34

Datatypes	34
Calling conventions	35
Building applications	35
Compilers	35
Linking	36
Include files	36
Using Microsoft C++	37
Using Borland C/C++	38
Setting up the Integrated Development Environment (IDE)	39
The module definition file	39
Using dynamic link libraries (DLLs)	40
Example programs	40

CHAPTER 4

Working with Databases

Connecting to databases	42
Creating database handles	42
Creating and populating a DPB	44
Adding parameters to a DPB	48
Attaching to a database	49
Requesting information about an attachment	51
Requesting buffer items and result buffer values	51
<i>isc_database_info()</i> call example	56
Disconnecting from databases	57
Deleting a database	58

CHAPTER 5

Working with Transactions

Starting transactions	60
Creating transaction handles	61
Creating a transaction parameter buffer	62
Calling <i>isc_start_transaction()</i>	70
Calling <i>isc_start_multiple()</i>	71
Ending transactions	73
Using <i>isc_commit_transaction()</i>	74

Using <i>isc_prepare_transaction2()</i>	77
Using <i>isc_rollback_transaction()</i>	77

CHAPTER 6 **Working with Dynamic SQL**

Overview of the DSQL programming process	79
DSQL API limitations	80
Accessing databases	80
Handling transactions	81
Creating a database	82
Processing Blob data	83
Processing array data	83
Writing an API application to process SQL statements	83
Determining if API calls can process an SQL statement	84
Representing an SQL statement as a character string	84
Specifying parameters in SQL statement strings	85
Understanding the <i>XSQLDA</i>	85
<i>XSQLDA</i> field descriptions	87
Input descriptors	89
Output descriptors	89
Using the <i>XSQLDA_LENGTH</i> macro	90
SQL datatype macro constants	90
Handling varying string datatypes	93
Handling NUMERIC and DECIMAL datatypes	93
Coercing datatypes	94
Aligning numerical data	95
DSQL programming methods	96
Method 1: Non-query statements without parameters	96
Method 2: Non-query statements with parameters	98
Method 3: Query statements without parameters	101
Method 4: Query statements with parameters	106
Determining an unknown statement type at runtime	113

CHAPTER 7 **Working with Blob Data**

What is a Blob?	118
----------------------------------	-----

	How are Blob data stored?	119
	Blob subtypes.	119
	Blob database storage	120
	Blob data operations	120
	Reading data from a Blob	121
	Writing data to a Blob	126
	Deleting a Blob	130
	Requesting information about an open Blob	131
	Item-list buffer items and result buffer values	131
	<i>isc_blob_info()</i> call example	133
	Blob descriptors	134
	Populating a Blob descriptor	135
	Filtering Blob data	136
	Using your own filters	136
	Declaring an external Blob filter to the database	137
	Writing an external Blob filter	137
	Writing an application that requests filtering.	142
CHAPTER 8	Working with Array Data	
	Introduction to arrays	150
	Array database storage	151
	Array descriptors.	151
	Populating an array descriptor	152
	Accessing array data	153
	Reading data from an array.	154
	Writing data to an array.	160
	Deleting an array.	166
CHAPTER 9	Working with Conversions	
	Converting date and times from InterBase to C format	168
	Converting dates from C to InterBase format	169
	Reversing byte order of numbers with <i>isc_vax_integer()</i>	170
CHAPTER 10	Handling Error Conditions	
	Setting up an error status vector	172

Using information in the status vector	172
Checking the status vector for errors	173
Displaying InterBase error messages	173
Capturing InterBase error messages	174
Setting an SQLCODE value on error	176
Displaying SQL error messages	176
Capturing SQL error messages	177
Parsing the status vector	178

CHAPTER 11 **Working with Events**

Understanding the event mechanism	188
Event parameter buffers	188
Synchronous event notification	189
Asynchronous event notification	189
Transaction control of events	189
Creating EPBs with <i>isc_event_block()</i>	190
Waiting on events with <i>isc_wait_for_event()</i>	191
Continuous processing with <i>isc_que_events()</i>	192
Creating an AST	193
A complete <i>isc_que_events()</i> example	193
Determining which events occurred with <i>isc_event_counts()</i>	196
Canceling interest in asynchronous events with <i>isc_cancel_events()</i>	197

CHAPTER 12 **Working with Services**

Overview of the Services API	200
General information	200
Using services parameter buffers	200
Attaching to the Services Manager with <i>isc_service_attach()</i>	202
Detaching from a Services Manager with <i>isc_service_detach()</i>	203
Invoking service tasks with <i>isc_service_start()</i>	204
Using request buffers	204
Overview of task identifiers	204

Backing up and restoring databases	205
Setting database properties	211
Invoking database maintenance	213
Requesting database and server status reports	216
Configuring users	217
Administering software activation certificates	219
Querying the Services Manager	220
Blocking and specifying timeout	220
Services API query example	221
Using result buffers	222
Querying server configuration	224
Querying security configuration	231
Querying service tasks	237
Using the Services API with Delphi and C++Builder	239

Part II: API Reference Guide

CHAPTER 13 **API Function Reference**

Function categories	243
Array functions	244
Blob functions	245
Database functions	246
Conversion functions	246
DSQL functions	247
Error-handling functions	248
Event functions	248
Information functions	249
Security functions	249
Services functions	250
Transaction control functions	250
Using function definitions	251
<i>isc_add_user()</i>	251
<i>isc_array_get_slice()</i>	254
<i>isc_array_lookup_bounds()</i>	259

<i>isc_array_lookup_desc()</i>	263
<i>isc_array_put_slice()</i>	266
<i>isc_array_set_desc()</i>	273
<i>isc_attach_database()</i>	276
<i>isc_blob_default_desc()</i>	279
<i>isc_blob_gen_bpb()</i>	281
<i>isc_blob_info()</i>	282
<i>isc_blob_lookup_desc()</i>	284
<i>isc_blob_set_desc()</i>	286
<i>isc_cancel_blob()</i>	287
<i>isc_cancel_events()</i>	289
<i>isc_close_blob()</i>	290
<i>isc_commit_retaining()</i>	291
<i>isc_commit_transaction()</i>	292
<i>isc_create_blob2()</i>	294
<i>isc_create_database()</i>	297
<i>isc_database_info()</i>	297
<i>isc_decode_sql_date()</i>	299
<i>isc_decode_sql_time()</i>	300
<i>isc_decode_timestamp()</i>	301
<i>isc_delete_user()</i>	302
<i>isc_detach_database()</i>	305
<i>isc_drop_database()</i>	306
<i>isc_dsql_allocate_statement()</i>	308
<i>isc_dsql_alloc_statement2()</i>	310
<i>isc_dsql_describe()</i>	312
<i>isc_dsql_describe_bind()</i>	314
<i>isc_dsql_execute()</i>	317
<i>isc_dsql_execute2()</i>	321
<i>isc_dsql_execute_immediate()</i>	325
<i>isc_dsql_exec_immed2()</i>	328
<i>isc_dsql_fetch()</i>	330
<i>isc_dsql_free_statement()</i>	334

<i>isc_dsqli_prepare()</i>	336
<i>isc_dsqli_set_cursor_name()</i>	339
<i>isc_dsqli_sql_info()</i>	342
<i>isc_encode_sql_date()</i>	344
<i>isc_encode_sql_time()</i>	345
<i>isc_encode_timestamp()</i>	346
<i>isc_event_block()</i>	347
<i>isc_event_counts()</i>	349
<i>isc_expand_dpb()</i>	351
<i>isc_get_segment()</i>	353
<i>isc_interprete()</i>	355
<i>isc_modify_user()</i>	356
<i>isc_open_blob2()</i>	360
<i>isc_prepare_transaction()</i>	362
<i>isc_prepare_transaction2()</i>	363
<i>isc_print_sqlerror()</i>	365
<i>isc_print_status()</i>	366
<i>isc_put_segment()</i>	367
<i>isc_que_events()</i>	369
<i>isc_rollback_retaining()</i>	373
<i>isc_rollback_transaction()</i>	375
<i>isc_service_attach()</i>	376
<i>isc_service_detach()</i>	377
<i>isc_service_query()</i>	378
<i>isc_service_start()</i>	380
<i>isc_sqlcode()</i>	381
<i>isc_sql_interprete()</i>	382
<i>isc_start_multiple()</i>	383
<i>isc_start_transaction()</i>	386
<i>isc_transaction_info()</i>	389
<i>isc_vax_integer()</i>	391
<i>isc_version()</i>	392
<i>isc_wait_for_event()</i>	394

APPENDIX A	InterBase Document Conventions	
	The InterBase documentation set	398
	Printing conventions	399
	Syntax conventions	400
APPENDIX B	Data Structures	
	Array descriptor	402
	Datatypes for array descriptors	403
	Blob descriptor	403
	Character sets	404
	Blob information buffers	404
	Blob buffer items	406
	Blob parameter buffer	406
	Database information request buffer and result buffer	407
	Request buffer	408
	Result buffer	408
	Request buffer items and result buffer values	409
	SQL datatype macro constants	418
	Status vector	420
	Meaning of the first long in a cluster	420
	Transaction parameter buffer	422
	XSQLDA and XSQLVAR	425
	XSQLDA field descriptions	427
	XSQLVAR field descriptions	428
	<i>Index</i>	<i>i</i>

List of Tables

Table 1.1	<i>API Guide</i> chapters	18
Table 3.1	Environment variables used by InterBase	33
Table 3.2	InterBase library file names	36
Table 3.3	Microsoft C compiler options	37
Table 3.4	Borland C compiler options	38
Table 4.1	API database functions	42
Table 4.2	DPB parameters	45
Table 4.3	Alphabetical list of DPB parameters	46
Table 4.4	DPB parameters recognized by <i>isc_expand_dpb()</i>	48
Table 4.5	<i>isc_expand_dpb()</i> parameters	48
Table 4.6	Database information items for database characteristics	52
Table 4.7	Database information items for environmental characteristics	54
Table 4.8	Database information items for performance statistics	55
Table 4.9	Database information items for operation counts	56
Table 5.1	API transaction functions	60
Table 5.2	Additional API transaction functions	60
Table 5.3	TPB constants	63
Table 5.4	Isolation level interaction with read and write operations	67
Table 6.1	SQL statements that cannot be processed by the API	84
Table 6.2	<i>XSQLDA</i> field descriptions	87
Table 6.3	<i>XSQLVAR</i> field descriptions	88
Table 6.4	SQL datatypes, macro expressions, and C datatypes	90
Table 6.5	SQL statement strings and recommended processing methods	96
Table 6.6	Statement types	114
Table 7.1	API Blob functions	118
Table 7.2	Blob request and return items	132
Table 7.3	Status message return items	132
Table 7.4	<i>isc_blob_ctl</i> structure field descriptions	140
Table 7.5	Action constants	142
Table 7.6	Blob parameter buffer parameter types	145
Table 8.1	API array access functions	150
Table 10.1	Error-handling functions	171
Table 10.2	Interpretation of status vector clusters	179

Table 10.3	<i>#defines</i> for status vector numeric descriptors	180
Table 11.1	API event functions	187
Table 12.1	Syntax of Services Manager connect string, by protocol	202
Table 12.2	Services API tasks	204
Table 12.3	Services API database backup arguments	206
Table 12.4	Services API database restore arguments	208
Table 12.5	Services API database properties arguments	211
Table 12.6	Services API database validation arguments.	213
Table 12.7	Services API database sweep arguments.	214
Table 12.8	Services API limbo transaction arguments	215
Table 12.9	Services API status report arguments	216
Table 12.10	Services API display users arguments	217
Table 12.11	Services API add user arguments	217
Table 12.12	Services API remove user arguments	218
Table 12.13	Services API software activation certificate arguments	219
Table 12.14	Services API server configuration query items	224
Table 12.15	Services API software activation certificate arguments	227
Table 12.16	Services API security configuration query items	231
Table 12.17	Services API user information arguments	232
Table 12.18	Services API database connection information arguments	235
Table 12.19	Services API task query items	237
Table 12.20	Services API limbo transaction arguments	238
Table 13.1	Array functions	244
Table 13.2	Blob functions	245
Table 13.3	Database functions	246
Table 13.4	Date and conversion functions	246
Table 13.5	DSQL functions	247
Table 13.6	Error-handling functions	248
Table 13.7	Event functions	248
Table 13.8	Information functions	249
Table 13.9	Security functions	249
Table 13.10	Service functions	250
Table 13.11	Transaction control functions	250
Table 13.12	Function description format	251
Table 13.13	Error messages for user security functions	253
Table 13.14	Datatypes for array descriptor fields	261

Table 13.15	Datatypes for array descriptor fields	264
Table 13.16	Datatypes for array descriptor fields	274
Table 13.17	Blob descriptor fields	280
Table 13.18	Blob descriptor fields	285
Table 13.19	Error messages for user security functions	304
Table 13.20	Error messages for user security functions	358
Table 13.21	Transaction information request item	389
Table 13.22	Status message return items	390
Table A.1	Books in the InterBase 6 documentation set	398
Table A.2	Text conventions	399
Table A.3	Syntax conventions	400
Table B.1	Array descriptor fields	402
Table B.2	Datatypes for array descriptors	403
Table B.3	Blob descriptor fields	404
Table B.4	Blob information items and return values	406
Table B.5	Status message return items	406
Table B.6	Blob parameter buffer parameter types	407
Table B.7	Status message return items	409
Table B.8	Database information items for database characteristics	410
Table B.9	Database information items for environmental characteristics	412
Table B.10	Database information items for performance statistics	413
Table B.11	Database information items for operation counts	414
Table B.12	DPB parameters	415
Table B.13	Alphabetical list of DPB parameters	416
Table B.14	SQL datatypes, macro expressions, and C datatypes	418
Table B.15	Interpretation of status vector clusters	421
Table B.16	<i>#defines</i> for status vector numeric descriptors	422
Table B.17	TPB constants	423
Table B.18	<i>XSQLDA</i> field descriptions	427
Table B.19	<i>XSQLVAR</i> field descriptions	428

Using the API Guide

The InterBase *API Guide* is a task-oriented explanation of how to write, preprocess, compile, and link database applications using the InterBase Applications Programming Interface (API), and a *host programming language*, either C or C++.

This chapter describes the focus of this book, and provides a brief overview of its chapters.

Who should use this guide

The InterBase *API Guide* is intended for knowledgeable database applications programmers. It assumes full knowledge of:

- SQL and dynamic SQL (DSQL).
- Relational database programming.
- C programming.

Topics covered in this guide

The *API Guide* is divided into two parts:

- A task-oriented user's guide that explains how to use API function calls to perform related database tasks, such as attaching to and detaching from a database.
- An API function call reference that describes the purpose of each function, its syntax, its parameters, requirements, restrictions, and return values, as well as examples of use and cross-references to related functions.

The following table provides a brief description of each chapter in the *API Guide*:

Chapter	Description
Chapter 2, "Application Requirements"	Describes support structures and elements common to programming with API calls
Chapter 3, "Programming with the InterBase API"	Describes special requirements for programming InterBase applications with the InterBase API
Chapter 4, "Working with Databases"	Describes how to attach to and detach from databases, and how to request information about attachments
Chapter 5, "Working with Transactions"	Explains how to start transactions in different modes, and how to commit them or roll them back
Chapter 6, "Working with Dynamic SQL"	Describes how to process DSQL data definition and data manipulation statements using API calls
Chapter 7, "Working with Blob Data"	Describes how to select, insert, update, and delete Blob data in applications
Chapter 8, "Working with Array Data"	Describes how to select, insert, update, and delete array data in applications
Chapter 9, "Working with Conversions"	Describes how to select, insert, update, and delete date and time data in applications, and how to reverse the byte order of numbers with <i>isc_vax_integer()</i>
Chapter 10, "Handling Error Conditions"	Describes how to trap and handle database errors in applications
Chapter 11, "Working with Events"	Explains how triggers interact with applications and describes how to register interest in events, wait on them, and respond to them in applications

TABLE 1.1 *API Guide* chapters

Chapter	Description
Chapter 12: “Working with Services”	
Chapter 13, “API Function Reference”	Describes the syntax of each function call in detail.
Appendix A, “InterBase Document Conventions”	Lists typefaces and special characters used in this book to describe syntax and identify object types.
Appendix B, “Data Structures”	Lists and describes the data structures, constants, and buffers that are defined in <code>ibase.h</code> .

TABLE 1.1 *API Guide* chapters (continued)

Sample database and applications

The InterBase **Examples** subdirectory contains a sample database and sample application source code. The examples in this *API Guide* make use of this sample database and source code wherever possible.

PART I

API User's Guide



Application Requirements

This chapter summarizes programming requirements for using categories of API functions in database applications, and provides cross-references to more detailed information in later chapters.

All API applications must use certain API functions and support structures. For example, all applications connect to at least one database, and run at least one transaction. All applications, therefore, must declare and initialize database handles and transaction handles. They may also need to declare and populate database parameter buffers (DPBs), transaction parameter buffers (TPBs), and service parameter buffers (SPBs). This chapter outlines those requirements, and points you to more detailed information later in this book.

Some API applications may use specific API functions, such as the functions that permit an application to process dynamic SQL (DSQL) statements. These applications have additional requirements that are also outlined in this chapter along with pointers to more detailed information elsewhere in this book.

Requirements for all applications

The following sections outline these requirements for all API applications:

- Including **ibase.h**
- Database requirements
- Transaction requirements

Including **ibase.h**

The InterBase subdirectory, **include**, contains the **ibase.h** header file, which should be included in all source code modules for API applications. **ibase.h** contains API function prototypes. It also contains structure typedefs, parameter definitions, and macros required by various API functions.

To include **ibase.h** in a source code module, insert the following `#include` near the start of the source code:

```
#include <ibase.h>
```

If **ibase.h** is not on your compiler's search path, you may need to provide a full path specification and enclose the file name in quotation marks.

Failure to include **ibase.h** can prevent the successful compilation and linking of an application.

Database requirements

All applications that work with databases must provide one database handle for each database to be accessed. A *database handle* is a long pointer that is used in API functions to attach to a database and to reference it in subsequent API calls. The InterBase header file, **ibase.h**, contains a *#define* useful for declaring database handles.

When establishing a connection to a database, optional database attachment characteristics, such as a user name and password combination, can be passed to the attachment through a database parameter buffer (DPB). Usually, one DPB is set up for each database attachment, although database attachments can also share a DPB.

► *Declaring database handles*

A database handle must be declared and initialized to zero before use. The following code illustrates how to declare and initialize a database handle:

```
#include <ibase.h>
. . .
/* Declare a database handle. */
isc_db_handle db1;
. . .
/* Initialize the handle. */
db1 = 0L;
```

For more information about declaring, initializing, and using database handles, see **Chapter 4, “Working with Databases.”**

► *Setting up a DPB*

A DPB is a byte array describing optional database attachment characteristics. A DPB must be set up and populated before attaching to a database. Parameters that can be passed to the DPB are defined in **ibase.h**.

For more information about setting up, populating, and using a DPB, see **Chapter 4, “Working with Databases.”**

Transaction requirements

All applications must provide one transaction handle for each transaction to be accessed. A *transaction handle* is a long pointer that is used in API functions to start a transaction and to reference it in subsequent API calls. The InterBase header file, **ibase.h**, contains a *#define* useful for declaring transaction handles.

When starting a transaction, optional transaction characteristics, such as access method and isolation level, can be passed to the start-up call through a transaction parameter buffer (TPB). Usually, one TPB is set up for each transaction, although transactions with the same operating characteristics can also share a TPB.

► *Declaring transaction handles*

A transaction handle must be declared and initialized to zero before use. The following code illustrates how to declare and initialize a transaction handle:

```
#include <ibase.h>
. . .
/* Declare a transaction handle. */
isc_tr_handle tr1;
```

```

. . .
/* Initialize the handle. */
tr1 = 0L;

```

For more information about declaring, initializing, and using transaction handles, see **Chapter 5, “Working with Transactions.”**

► *Setting up a TPB*

A TPB is a byte array containing parameters that describe optional transaction characteristics. In these cases, the TPB must be set up and populated before starting a transaction. Parameters that can be passed to the TPB are defined in **ibase.h**.

For more information about setting up, populating, and using a TPB, see **Chapter 5, “Working with Transactions.”**

Additional requirements

The following sections outline possible additional requirements for API applications developed on certain system platforms, such as Microsoft Windows, and for general classes of API functions, such as those that process DSQL statements.

Microsoft Windows requirements

InterBase client applications for Microsoft Windows have programming requirements specific to that environment and the C/C++ compilers available there.

The InterBase header file, **ibase.h**, provides prototypes of all API functions. For Windows applications, these prototypes make use of the following declarations:

```

#define ISC_FAR __far
#define ISC_EXPORT ISC_FAR __cdecl __loadds __export

```

For example, the `isc_attach_database()` prototype in **ibase.h** is:

```

ISC_STATUS ISC_EXPORT isc_attach_database(
    ISC_STATUS ISC_FAR *,
    short,
    char ISC_FAR,
    isc_db_handle ISC_FAR *,
    short,
    char ISC_FAR *);

```

When Windows client applications make calls and cast C datatypes, they should make explicit use of the *ISC_FAR* declaration.

Note The *ISC_EXPORT* keyword is omitted from the API function reference because on all non-Windows platforms it is undefined.

For more information about Windows requirements, see **Chapter 3, “Programming with the InterBase API.”**

DSQL requirements

API applications that build or prompt for DSQL queries at run time require careful declaration, initialization, and population of extended SQL descriptor area (*XSQLDA*) structures for data transfer to and from the database. In addition, many API functions, such as *isc_dsql_allocate_statement()* and *isc_dsql_describe()*, also make use of statement handles for DSQL processing.

ibase.h provides typedefs for the *XSQLDA* structure, and its underlying structure, the *XSQLVAR*. It also provides a *#define* for the statement handle, a macro for allocating the appropriate amount of space for an instance of an *XSQLDA* in an application, and *#defines* for DSQL information parameters passed to *isc_dsql_sql_info()*.

The following code illustrates how to declare an *XSQLDA* structure for use in an application, and how to declare a statement handle:

```
#include <ibase.h>
. . .
XSQLDA *insqlda;
isc_stmt_handle sql_stmt;
. . .
```

For more information about DSQL programming with the API, see **Chapter 6, “Working with Dynamic SQL.”**

Blob requirements

To work with Blob data that must be filtered, an API application must set up a Blob parameter buffer (BPB) for each Blob. A BPB is a variable-length byte vector declared in an application to store control information that dictates Blob access. The BPB can contain a number of constants, defined in **ibase.h**, that describe the Blob and the Blob subtypes that specify Blob filtering.

IMPORTANT Blob filtering is not available on NetWare servers.

Applications that work with Blob data in an international environment must also declare and populate a Blob descriptor that contains character set information for the Blob. The Blob descriptor structure is defined in **ibase.h**. To declare a Blob descriptor, an application must provide code like this:

```
#include <ibase.h>
. . .
ISC_BLOB_DESC to_desc;
```

Except on NetWare servers, where they are not supported, Blob filters enable a Blob to be translated from one format to another, such as from a compressed state to an decompressed state or vice versa. If Blob filters are desired, separate filter functions must be created and defined to the database to ensure their use when Blob data is accessed.

Finally, to access Blob data, applications must make extensive use of API DSQL functions.

For more information about working with Blob data and Blob filters, see **Chapter 7, “Working with Blob Data.”** For more information about DSQL, see **Chapter 6, “Working with Dynamic SQL.”**

Array requirements

API functions that handle array processing require the use of an array descriptor structure and array IDs, defined in **ibase.h**. In addition, applications accessing arrays must make extensive use of API DSQL functions.

The following code illustrates how to declare an array descriptor and array ID variable, and how to initialize an array ID to zero before use:

```
#include <ibase.h>
. . .
ISC_ARRAY_DESC desc;
ISC_QUAD array_id;
. . .
array_id = 0L;
. . .
```

For more information about working with arrays, see **Chapter 8, “Working with Array Data.”** For more information about DSQL, see **Chapter 6, “Working with Dynamic SQL.”**

Event requirements

InterBase events are messages passed from a trigger or stored procedure to an application to announce the occurrence of specified conditions or actions, usually database changes such as insertions, modifications, or deletions of records.

Before an application can respond to an event, it must register interest in an event. To register interest in an event, the application must establish and populate two event parameter buffers (EPBs), one for holding the initial occurrence count values for each event of interest, and another for holding the changed occurrence count values. These buffers are passed as parameters to several API event functions, and are used to determine which events have occurred.

In C, each EPB is declared as a char pointer, as follows:

```
char *event_buffer, *result_buffer;
```

Once the buffers are declared, *isc_event_block()* is called to allocate space for them, and to populate them with starting values.

For more information about events, see **Chapter 11, “Working with Events.”**

Error-handling requirements

Most API functions return status information in an *error status vector*, an array of 20 longs. To handle InterBase error conditions, should they arise, applications should declare a status vector as follows:

```
#include <ibase.h>
. . .
ISC_STATUS status_vector[20];
```

ISC_STATUS is a *#define* in **ibase.h** provided for programming convenience and platform independence.

ibase.h also contains *#defines* for all InterBase error conditions. Applications can use API error-handling functions to construct error messages from the status vector that are based on these error conditions, or can examine the status vector directly for particular error conditions using the *#defines* in place of error numbers. Using *#defines* in this manner makes source code easier to understand and maintain.

For more information about error handling, see **Chapter 11, “Working with Events.”**

Services requirements

InterBase provides an API interface to enable your applications to request information about server and database properties, and to invoke tasks to administer servers and databases. Your application can initiate a connection to a local instance of the InterBase server, or to a remote server over a network. Through this connection, your application submits requests to the server and receives resultant data.

For more information about using this API facility, see **Chapter 12, “Working with Services.”**

Compiling and linking

On most development platforms, an API application is compiled like any standard C or C++ application. For more information about a particular compiler, consult the compiler’s documentation.

On most platforms, InterBase supports dynamic linking of its library at run time. One exception to this scenario is on Microsoft Windows, where an application must explicitly link to the InterBase library (`gds32.lib` or `gds32_ms.lib`).

On Microsoft Windows, there are particular compiling options to be aware of. For more information about linking under Windows, see **Chapter 3, “Programming with the InterBase API.”**

For all other platforms, see the InterBase *Embedded SQL Guide* for specific compiling and linking guidelines.

Programming with the InterBase API

This chapter provides information specific to programming InterBase applications on a client with C/C++. It assumes familiarity with Borland C/C++ or Microsoft C/C++, InterBase, and the InterBase documentation set, particularly the *Language Reference*.

Basic procedure for application development

The basic steps in application development using the InterBase Client are:

- Determine which client and server platforms the application will run on. InterBase clients and servers include Microsoft Windows 95/98, Windows NT, Linux, and several brands of UNIX. The InterBase server also runs on Novell NetWare 4.
- Code the application in C or C++.
- Compile and link the application.
- Test and debug the application.
- Deploy the application on the production client platform.

Supported development environments

The InterBase client library enables developers to design InterBase SQL client applications that connect to remote InterBase servers on Windows 95/98, Windows NT, UNIX, or NetWare.

See the *Operations Guide* for more specific information about this topic.

User name and password requirements

When an InterBase client application is compiled, linked, and run, the client must *always* send a valid user name and password combination to the InterBase server. The server checks the user name and password against the user name and password combinations stored in its security database. If a match is found, the client can attach to InterBase databases on the server. If a match is not found, the server denies the attachment request.

For a successful attachment to occur, the following steps must be taken:

1. A user with SYSDBA privileges must add a client's user name and password to the server's security database (**isc4.gdb**). Use the Server Manager to do this on Windows platforms. On UNIX, use the **gsec** utility.
2. The client must send a valid user name and password combination to the server. Password is case-sensitive.

Note Under some circumstances, you can connect to a database even if you don't have a user name in the InterBase security database. In order for this to happen, the following things must be true:

- Both the client and server are running under UNIX
- Your current login exists on the server host
- You are logging in from a trusted client; a trusted client is one that is listed in the **/etc/hosts.equiv** or **/etc/gds_hosts.equiv** file on the server or in the **.rhosts** file in your home directory on the server
- You have *not* specified a user name and password in the connect string

Note InterBase comes with one user ID predefined. This is the SYSDBA user ID. The default password is **masterkey**. This user ID is for use by the database administrator, and it has special privileges that are not available to any other user ID. Do not use this user ID for end-users' application.

Specifying user name and password

A client application must specify a user name and password when it attaches to a database. Failure to provide a valid user name and password combination results in an error. Use the following methods to provide user names and passwords:

- Create a database parameter block (DPB) with *isc_dpb_user_name* and *isc_dpb_password*, and pass the parameter block using *isc_attach_database()*.
- Add *isc_dpb_user_name* and *isc_dpb_password* parameters to an existing DPB with *isc_expand_dpb()*.

For more information about the DPB, *isc_attach_database()*, and *isc_expand_dpb()*, see **Chapter 4, “Working with Databases.”**

Using environment variables

InterBase client applications can use three environment variables to establish program parameters. These variables must be set so that they are available to the application when it is running. For example, setting these variables within a DOS window after Windows has been started does not affect any Windows programs, but affects DOS applications in that window.

The following table summarizes these variables and their uses:

Variable	Purpose	Example
<i>ISC_DATABASE</i>	Specifies a default server and database directory to use on the remote server	<code>SET ISC_DATABASE = ingold:/usr/interbase/examples</code>
<i>ISC_USER</i>	Specifies a user name for the PC client application	<code>SET ISC_USER = HERMES</code>
<i>ISC_PASSWORD</i>	Specifies a case-sensitive password for the PC client application	<code>SET ISC_PASSWORD = Ichneumon</code>

TABLE 3.1 Environment variables used by InterBase

The *ISC_USER* and *ISC_PASSWORD* environment variables are used together to establish a valid user name and password combination to pass to the remote InterBase database server.

IMPORTANT Do not use the `ISC_PASSWORD` environment variable when security is a concern. Anyone with access to a client where an `ISC_PASSWORD` environment variable is defined in a file such as `autoexec.bat` can easily view the password.

Setting a default database directory

To connect automatically to a default database directory on a remote server, create the `ISC_DATABASE` environment variable and set it to the full path specification for the desired database directory, including host and path names.

Note Host name specification is specific to the server's operating system and network protocol. The host syntax in the previous example is for a generic UNIX server. For other servers and operating systems, see that system's reference manuals.

Setting a user name and password

To set up a default user name and password for use on a PC client, create two environment variables, `ISC_USER`, and `ISC_PASSWORD`.

Even if `ISC_USER` and `ISC_PASSWORD` are set, a different user name and password may be specified in a DPB used as an argument to `isc_attach_database()`. A user name or password specified in a database parameter block overrides the OS environment variables.

Note Using environment variables in this manner is not secure, and therefore not recommended.

Datatypes

InterBase supports a wide variety of datatypes for application development. These datatypes are defined in a typedef to be platform-independent. The InterBase client libraries are also compiled with packed data structures to be compatible with a variety of platforms.

For more information about InterBase datatypes, see the *Language Reference*.

Calling conventions

Conventions for calling functions vary from platform to platform. Specifically:

- On UNIX platforms, use the C calling conventions (cdecl) in all cases.
- On Windows 95 and Windows NT, use the standard calling conventions (_stdcall) for all functions that have a fixed number of arguments. There are only three functions that have a variable number of arguments. For these three—*isc_start_transaction()*, *isc_expand_dpb()*, and *isc_event_block()*—use the cdecl conventions.

Building applications

This section discusses compilers and libraries that are needed to build InterBase applications.

HELP WITH LINKING AND COMPILING On each platform, there is a **makefile** in the **examples** directory that contains detailed platform-specific information about linking and compiling. Open the makefile in a text editor to access the information.

Compilers

The import libraries included with InterBase have been tested with the following compilers:

Windows platforms

- Borland C++ 5.0
- Microsoft Visual C++ 2.0
- Microsoft Visual C++ 4.0

Solaris

- C SPARCWorks SC4.2 C compiler
- C++ SPARCWorks SC3.0.1 C++ compiler
- COBOL MicroFocus Cobol 4.0
- ADA SPARCWorks SC4.0 Ada compiler
- FORTRAN SPARCWorks SC4.0 Fortran compiler

HP-UX

- C HP C/HP-UX Version A.10.32
- C++ HP C++/HP-UX Version A.10.22
- COBOL MicroFocus Cobol 4.0
- ADA Alsys Ada - AdaWorld V5.5.4
- FORTRAN HP Fortran/9000 10.20 Release

Linking

The InterBase library files reside in the **lib** subdirectory of the installation directory. Applications must link with the InterBase client library. This library name varies depending on the platform and the compiler.

Platform/compiler	InterBase library file
Windows/Borland C++	gds32.lib
Windows/Microsoft Visual C++ 2.0 and 4.0	gds32_ms.lib
Solaris/all	gdsmt
HPUX/all	gds

TABLE 3.2 InterBase library file names

Borland compilers earlier than 5.0 do not work with **gds32.lib**.

Include files

Applications must include the **ibase.h** header file to pick up the InterBase type definitions and function prototypes. This file is in the **include** subdirectory of the InterBase install directory.

On UNIX platforms, the **gds.h** file is available in the installation directory for backward compatibility.

Using Microsoft C++

Use the following options when compiling applications with Microsoft C++:

Option	Action
<code>c</code>	Compile without linking (DLLs only)
<code>Zi</code>	Generate complete debugging information
<code>DWIN32</code>	Defines "WIN32" to be the null string
<code>D_MT</code>	Use a multi-thread, statically-linked library

TABLE 3.3 Microsoft C compiler options

For example, these commands use the Microsoft compiler to build a DLL that uses InterBase:

```
cl -c -Zi -DWIN32 -D_MT -LD udf.c
    lib -out:udf.lib -def:funclib.def -machine:i586 -subsystem:console
    link -DLL -out:funclib.dll -DEBUG:full,mapped -DEBUGTYPE:CV
    -machine:i586 -entry:_DllMainCRTStartup@12 -subsystem:console
    -verbose udf.obj udf.exp gds32.lib ib_util_ms.lib crtdll.lib
```

This command builds an InterBase executable using the Microsoft compiler:

```
cl -Zi -DWIN32 -D_MT -MD udfctest.c udf.lib gds32.lib
    ib_util_ms.lib crtdll.lib
```

Note See “Working with UDFs” in the *Developer’s Guide* for more about compiling and linking user-defined libraries.

Using the Dynamic Runtime Library If you are

- using a Microsoft Visual C++ 2.0 or Microsoft Visual C++ 4.0
- compiling and linking separately, and
- using the Dynamic Runtime Library (`msvcrt20.dll` or `msvcrt40.dll`)

you need to use the `/MD` compiler flag to compile with the run time library (RTL), as well as linking with the correct import library.

Using Borland C/C++

Use the following options when compiling applications with Borland C++:

Option	Action
<i>v</i>	Turns on source debugging
<i>a4</i>	Structure padding/byte alignment
<i>DWIN32</i>	<ul style="list-style-type: none"> • Defines the string "WIN32" • With no argument, it defines it to the null string
<i>tWM</i>	Makes the target multi-threaded
<i>tWC:</i>	<ul style="list-style-type: none"> • Makes the target a console .EXE with all functions exportable • Cannot be used with the <i>-tWCD</i> option
<i>tWCD</i>	Makes the target a console .DLL with all functions exportable; cannot be used with the <i>-tWC</i> option

TABLE 3.4 Borland C compiler options

The following command creates a DLL named **funclib.dll** from a source file named **udf.c**:

```
implib mygds32.lib \interbas\bin\gds32.dll
bcc32 -v -a4 -DWIN32 -tWM -tWCD -efunclib.dll udf.c mygds32.lib
```

The following commands create an InterBase executable named **udftest.exe** (which calls into **funclib.dll**) from a source file named **udftest.e** containing embedded SQL commands.

```
implib udf.lib funclib.dll
gpre -e udfest.e
bcc32 -v -a4 -DWIN32 -tWM -tWC udfest.c udf.lib mygds32.lib
```

When linking applications with Borland C command line linker, use the */c* option (case sensitive link).

Note There are equivalent general linker options within the Borland Integrated Development Environment (IDE). The default in the IDE is case-sensitive link (*/c* option) alone, which causes unresolved linker errors for all of the InterBase entry points.

Setting up the Integrated Development Environment (IDE)

The Borland Integrated Development Environment (IDE) offers options that are equivalent to the command line options.

► *IDE default*

The case-sensitive link (/c option) is the default in the IDE.

► *IDE Project Options dialog box*

Choose the following options from the IDE Project Options dialog box. The corresponding command-line option is also listed.

DIRECTORIES

Include directory: *interbase_home_dir*\include

Library directory: *interbase_home_dir*\lib

Note The default InterBase home directory is `c:\Program Files\InterBase Corp\InterBase`.

COMPILER

Source language compliance: Borland extensions

32-bit Compiler

Data alignment: Byte (-a4 option for 4 byte alignment)

LINKER

Choose Case-sensitive link ON (/c option).

The module definition file

Creating a module definition file can solve certain issues that arise during linking and compiling with the Borland C++ Builder:

- Set the STACKSIZE parameter to at least 10 kilobytes (10,240 bytes); 16 kilobytes (16,384 bytes) is recommended. A sample .def file is included in the **examples** subdirectory of the InterBase installation directory.

- Because the Borland C++Builder prepends an underscore to some API functions that **gds32.dll** exports without the underscore, you may need to add aliases for these functions to your module definition file, as in the following example:

```
IMPORTS
    _isc_start_transaction = GDS32.isc_start_transaction
```

Using dynamic link libraries (DLLs)

InterBase applications use the **gds32.dll** dynamic link library, which in turn loads the appropriate network DLLs. These DLLs unload automatically when the last calling application terminates. If the calling application exits abnormally (for example, from a protection fault), it is possible that DLLs will not be unloaded from memory. If this occurs, exit and restart Windows to free the resources.

Example programs

Example programs demonstrating how to use the InterBase API are included in the **examples** subdirectory of the InterBase installation directory. There is also a sample **.def** file.

On NT, there are two make files, **makefile.bc** for the Borland compiler and linker, and **makefile.msc** for the Microsoft compiler and linker. In both files, you must modify the **IBASE** environment variable to point to an absolute path.

In the **.bc** make file, modify the **BCDIR** variable to point to the absolute path to the Borland compiler and linker.

In the **.msc** make file, modify the **MSCDIR** variable to point to the absolute path to the Microsoft compiler and linker.

To build the example applications on NT using Borland C++, use the following command:

```
make -B -f makefile.bc all
```

To build the example applications using Microsoft C++, use this command:

```
nmake -B -f makefile.msc all
```

On UNIX systems, the command to build the example applications is as follows:

```
make all
```

Working with Databases

This chapter describes how to set up a *database parameter buffer* (DPB) that specifies database attachment parameters, how to set up and initialize database handles, and how to use the five API functions that control database access. It also explains how to set up item request and return buffers prior to retrieving information about an attached database.

The following table lists the API functions for working with databases. The functions are listed in the order that they typically appear in an application.

Call	Purpose
<i>isc_expand_dpb()</i>	Specifies additional parameters for database access, such as user names and passwords elicited from a user at run time; uses a previously declared and populated DPB
<i>isc_attach_database()</i>	Connects to a database and establishes initial parameters for database access, such as number of cache buffers to use; uses a previously declared and populated DPB
<i>isc_database_info()</i>	Retrieves requested information about an attached database, such as the version of the on-disk structure (ODS) that it uses
<i>isc_detach_database()</i>	Disconnects from an attached database and frees system resources allocated to that attachment
<i>isc_drop_database()</i>	Deletes a database and any support files, such as shadow files

TABLE 4.1 API database functions

Connecting to databases

Connecting to one or more databases is a four-step process:

1. Creating and initializing a database handle for each database to be attached.
2. Creating and populating a DPB for each database to be attached.
3. Optionally calling *isc_expand_dpb()* prior to actual attachment to add more database parameters to a previously created and populated DPB.
4. Calling *isc_attach_database()* for each database to which to connect.

These steps are described in the following sections of this chapter.

Creating database handles

Every database that is accessed in an application must be associated with its own *database handle*, a pointer to a FILE structure that is used by all API database functions. The **ibase.h** header file contains the following C typedef declaration for database handles:

```
typedef void ISC_FAR *isc_db_handle;
```

To use this typedef for declaring database handles in an application, include **ibase.h** in each source file module:

```
#include <ibase.h>
```

► *Declaring database handles*

To establish database handles for use, declare a variable of type *isc_db_handle* for each database that will be accessed at the same time. The following code declares two handles:

```
#include <ibase.h>
. . .
isc_db_handle db1;
isc_db_handle db2;
```

Once a database is no longer attached, its handle can be assigned to a different database in a subsequent attachment. If an application accesses several databases, but only accesses a subset of databases at the same time, it is only necessary to declare as many handles as there will be simultaneous database accesses. For example, if an application accesses a total of three databases, but only attaches to two of them at a time, only two database handles need be declared.

► *Initializing database handles*

Before a database handle can be used to attach to a database, it must be set to zero. The following code illustrates how two database handles are set to zero:

```
#include <ibase.h>
. . .
isc_db_handle db1;
isc_db_handle db2;
. . .
/* Set database handles to zero before attaching to a database. */
db1 = 0L;
db2 = 0L;
```

Once a database handle is initialized to zero, it can be used in a call to *isc_attach_database()* to establish a database connection. If a nonzero database handle is passed to *isc_attach_database()*, the connection fails and an error code is returned. For more information about establishing a database connection with *isc_attach_database()*, see **“Attaching to a database” on page 49**.

Creating and populating a DPB

Database attachments can optionally be tailored in many ways by creating a database parameter buffer (DPB), populating it with desired database characteristics, and passing the address of the DPB to `isc_attach_database()`.

For example, the DPB can contain a user name and password for attaching to a database on a remote server, and it might also contain a parameter that activates a database shadow file. For a list of all possible DPB parameters, see **Table 4.2, “DPB parameters,” on page 45**.

Usually a separate DPB is created for each database attachment, but if different attachments use the same set of parameters, they can share a DPB. If a DPB is not created or is not passed to `isc_attach_database()`, the database attachment uses a default set of parameters.

Tip Some of the DPB parameters correspond directly to **gfix** options. In fact, that’s how **gfix** is implemented: it sets certain DPB parameters and attaches to a database, where it performs the requested operation (sweep, set async writes, shutdown, and so on).

A DPB is a *char* array variable declared in an application, that consists of the following parts:

1. A byte specifying the version of the parameter buffer, always the compile-time constant, `isc_dpb_version1`.
2. A contiguous series of one or more *clusters* of bytes, each describing a single parameter.

Each cluster consists of the following parts:

1. A one-byte parameter type. There are compile-time constants defined for all the parameter types (for example, `isc_dpb_num_buffers`).
2. A one-byte number specifying the number of bytes that follow in the remainder of the cluster.
3. A variable number of bytes, whose interpretation (for example, as a number or as a string of characters) depends on the parameter type.

For example, the following code creates a DPB with a single parameter that sets the number of cache buffers to use when connecting to a database:

```
char dpb_buffer[256], *dpb, *p;
short dpb_length;
/* Construct the database parameter buffer. */
dpb = dpb_buffer;
*dpb++ = isc_dpb_version1;
```

```
*dpb++ = isc_num_buffers;
*dpb++ = 1;
*dpb++ = 90;
dpb_length = dpb - dpb_buffer;
```

IMPORTANT All numbers in the database parameter buffer must be represented in a generic format, with the least significant byte first, and the most significant byte last. Signed numbers should have the sign in the last byte. The API function *isc_vax_integer()* can be used to reverse the byte order of a number. For more information, see **“isc_vax_integer()” on page 391**.

The following table lists DPB items by purpose:

User validation	
User name	<i>isc_dpb_user_name</i>
Password	<i>isc_dpb_password</i>
Encrypted password	<i>isc_dpb_password_enc</i>
Role name	<i>isc_dpb_sql_role_name</i>
System database administrator’s user name	<i>isc_dpb_sys_user_name</i>
Authorization key for a software license	<i>isc_dpb_license</i>
Database encryption key	<i>isc_dpb_encrypt_key</i>
Environmental control	
Number of cache buffers	<i>isc_dpb_num_buffers</i>
<i>dbkey</i> context scope	<i>isc_dpb_dbkey_scope</i>
System management	
Force writes to the database to be done asynchronously or synchronously	<i>isc_dpb_force_write</i>
Specify whether or not to reserve a small amount of space on each database page for holding backup versions of records when modifications are made	<i>isc_dpb_no_reserve</i>
System management	
Specify whether or not the database should be marked as damaged	<i>isc_dpb_damaged</i>
Perform consistency checking of internal structures	<i>isc_dpb_verify</i>

TABLE 4.2 DPB parameters

Shadow control	
Activate the database shadow, an optional, duplicate, in-sync copy of the database	<i>isc_dpb_activate_shadow</i>
Delete the database shadow	<i>isc_dpb_delete_shadow</i>
Replay logging system control	
Activate a replay logging system to keep track of all database calls	<i>isc_dpb_begin_log</i>
Deactivate the replay logging system	<i>isc_dpb_quit_log</i>
Character set and message file specification	
Language-specific message file	<i>isc_dpb_lc_messages</i>
Character set to be utilized	<i>isc_dpb_lc_ctype</i>

TABLE 4.2 DPB parameters (*continued*)

The following table lists DPB parameters in alphabetical order. For each parameter, it lists its purpose, the length, in bytes, of any values passed with the parameter, and the value to pass.

Parameter	Purpose	Length	Value
<i>isc_dpb_activate_shadow</i>	Directive to activate the database shadow, which is an optional, duplicate, in-sync copy of the database	1 (Ignored)	0 (Ignored)
<i>isc_dpb_damaged</i>	Number signifying whether or not the database should be marked as damaged 1 = mark as damaged 0 = do <i>not</i> mark as damaged	1	0 or 1
<i>isc_dpb_dbkey_scope</i>	Scope of <i>dbkey</i> context. 0 limits scope to the current transaction, 1 extends scope to the database session	1	0 or 1
<i>isc_dpb_delete_shadow</i>	Directive to delete a database shadow that is no longer needed	1 (Ignored)	0 (Ignored)
<i>isc_dpb_encrypt_key</i>	String encryption key, up to 255 characters	Number of bytes in string	String containing key

TABLE 4.3 Alphabetical list of DPB parameters

Parameter	Purpose	Length	Value
<i>isc_dpb_force_write</i>	Specifies whether database writes are synchronous or asynchronous. 0 = asynchronous; 1 = synchronous	1	0 or 1
<i>isc_dpb_lc_ctype</i>	String specifying the character set to be utilized	Number of bytes in string	String containing character set name
<i>isc_dpb_lc_messages</i>	String specifying a language-specific message file	Number of bytes in string	String containing message file name
<i>isc_dpb_license</i>	String authorization key for a software license	Number of bytes in string	String containing key
<i>isc_dpb_no_reserve</i>	Specifies whether or not a small amount of space on 1 each database page is reserved for holding backup versions of records when modifications are made; keep backup versions on the same page as the primary record to optimize update activity 0 (default) = reserve space 1 = do not reserve space		0 or 1
<i>isc_dpb_num_buffers</i>	Number of database cache buffers to allocate for use with the database; default=75		Number of buffers to allocate
<i>isc_dpb_password</i>	String password, up to 255 characters	Number of bytes in string	String containing password
<i>isc_dpb_password_enc</i>	String encrypted password, up to 255 characters	Number of bytes in string	String containing password
<i>isc_dpb_sys_user_name</i>	String system DBA name, up to 255 characters	Number of bytes in string	String containing SYSDBA name
<i>isc_dpb_user_name</i>	String user name, up to 255 characters	Number of bytes in string	String containing user name

TABLE 4.3 Alphabetical list of DPB parameters (continued)

Note Some parameters, such as *isc_dpb_delete_shadow*, are directives that do not require additional parameters. Even so, you *must* still provide length and value bytes for these parameters. Set length to 1 and value to 0. InterBase ignores these parameter values, but they are required to maintain the format of the DPB.

Adding parameters to a DPB

Sometimes it is useful to add parameters to an existing DPB at run time. For example, when an application runs, it might determine a user's name and password and supply those values dynamically. The *isc_expand_dpb()* function can be used to pass the following additional parameters to a previously created and populated DPB at run time:

Parameter	Purpose
<i>isc_dpb_user_name</i>	String user name, up to 255 characters
<i>isc_dpb_password</i>	String password, up to 255 characters
<i>isc_dpb_lc_messages</i>	String specifying a language-specific message file
<i>isc_dpb_lc_ctype</i>	String specifying the character set to be utilized

TABLE 4.4 DPB parameters recognized by *isc_expand_dpb()*

IMPORTANT If you expect to add any of these parameters at run time, then create a larger than necessary DPB before calling *isc_expand_dpb()*, so that this function does not need to reallocate DPB storage space at run time. *isc_expand_dpb()* can reallocate space, but that space is not automatically freed when the database is detached.

isc_expand_dpb() requires the following parameters:

Parameter	Type	Description
<i>dpb</i>	<i>char **</i>	Pointer to a DPB
<i>dpb_size</i>	<i>unsigned short *</i>	Pointer to the current size, in bytes, of the DPB
...	<i>char *</i>	Pointers to item type and items to add to the DPB

TABLE 4.5 *isc_expand_dpb()* parameters

The third parameter in the table, "...", indicates a variable number of replaceable parameters, each with different names, but each a character pointer.

The following code demonstrates how *isc_expand_dpb()* is called to add a user name and password to the DPB after they are elicited from a user at run time:

```
char dpb_buffer[256], *dpb, *p;
char uname[256], upass[256];
short dpb_length;

/* Construct a database parameter buffer. */
```

```

dpb = dpb_buffer;
*dpb++ = isc_dpb_version1;
*dpb++ = isc_num_buffers;
*dpb++ = 1;
*dpb++ = 90;
dpb_length = dpb - dpb_buffer;
/* Now ask user for name and password. */
prompt_user("Enter your user name: ");
gets(uname);
prompt_user("\nEnter your password: ");
gets(upass);
/* Add user name and password to DPB. */
dpb = dpb_buffer;
isc_expand_dpb(&dpb, &dpb_length,
    isc_dpb_user_name, uname,
    isc_dpb_password, upass,
    NULL);

```

Attaching to a database

After creating and initializing a database handle, and optionally setting up a DPB to specify connection parameters, use *isc_attach_database()* to establish a connection to an existing database. Besides allocating system resources for the database connection, *isc_attach_database()* also associates a specific database with a database handle for use in subsequent API calls that require a handle.

isc_attach_database() expects six parameters:

- A pointer to an error status array, where attachment errors can be reported should they occur.
- The length, in bytes, of the database name for the database to open. If the database name includes a node name and path, these elements must be counted in the length argument.
- A string containing the name of the database to attach. The name can include a node name and path specification.
- A pointer to a previously declared and initialized database handle with which to associate the database to attach. All subsequent API calls use the handle to specify access to this database.
- The length, in bytes, of the DPB. If no DPB is passed, set this value to zero.
- A pointer to the DPB. If no DPB is passed, set this to NULL.

Each database attachment requires a separate call to *isc_attach_database()*.

The following code establishes an attachment to the InterBase example database, **employee.gdb**, and specifies a DPB to use for the attachment:

```
#include <ibase.h>
. . .
isc_db_handle db1;
char dpb_buffer[256], *dpb, *p;
short dpb_length;
char *str = "employee.gdb";
ISC_STATUS status_vector[20];
. . .
/* Set database handle to zero before attaching to a database. */
db1 = 0L;
/* Initialize the DPB. */
dpb = dpb_buffer;
*dpb++ = isc_dpb_version1;
*dpb++ = isc_num_buffers;
*dpb++ = 1;
*dpb++ = 90;
dpb_length = dpb - dpb_buffer;
/* Attach to the database. */
isc_attach_database(status_vector, strlen(str), str, &db1,
    dpb_length,
    dbp_buffer);
if (status_vector[0] == 1 && status_vector[1])
{
    error_exit();
}
. . .
```

The following code illustrates how to attach to a database without passing a DPB:

```
#include <ibase.h>
. . .
isc_db_handle db1;
char *str = "employee.gdb";
ISC_STATUS status_vector[20];
. . .
/* Set database handle to zero before attaching to a database. */
db1 = 0L;
/* Attach to the database. */
isc_attach_database(status_vector, strlen(str), str, &db1, 0, NULL);
if (status_vector[0] == 1 && status_vector[1])
```

```

{
    error_exit();
}
. . .

```

Requesting information about an attachment

After an application attaches to a database, it may need information about the attachment. The `isc_database_info()` call enables an application to query for attachment information, such as the version of the on-disk structure (ODS) used by the attachment, the number of database cache buffers allocated, the number of databases pages read from or written to, or write-ahead log information.

In addition to a pointer to the error status vector and a database handle, `isc_database_info()` requires two application-provided buffers, a request buffer, where the application specifies the information it needs, and a result buffer, where InterBase returns the requested information. An application populates the request buffer with information prior to calling `isc_database_info()`, and passes it both a pointer to the request buffer, and the size, in bytes, of that buffer.

The application must also create a result buffer large enough to hold the information returned by InterBase. It passes both a pointer to the result buffer, and the size, in bytes, of that buffer, to `isc_database_info()`. If InterBase attempts to pass back more information than can fit in the result buffer, it puts the value, `isc_info_truncated`, defined in `ibase.h`, in the final byte of the result buffer.

Requesting buffer items and result buffer values

The request buffer is a char array into which is placed a sequence of byte values, one per requested item of information. Each byte is an *item type*, specifying the kind of information desired. Compile-time constants for all item types are defined in `ibase.h`.

The result buffer returns a series of clusters of information, one per item requested. Each cluster consists of three parts:

1. A one-byte *item return type*. There are compile-time constants defined for all the item return types in `ibase.h`.
2. A two-byte number specifying the number of bytes that follow in the remainder of the cluster.

3. A *value*, stored in a variable number of bytes, whose interpretation (for example, as a number or as a string of characters) depends on the item return type.

A calling program is responsible for interpreting the contents of the result buffer and for deciphering each cluster as appropriate. In many cases, the value simply contains a number or a string (sequence of characters). But in other cases, the value is a number of bytes whose interpretation depends on the item return type.

The clusters returned to the result buffer are not aligned. Furthermore, all numbers are represented in a generic format, with the least significant byte first, and the most significant byte last. Signed numbers have the sign in the last byte. Convert the numbers to a datatype native to your system, if necessary, before interpreting them. The API call, *isc_vax_integer()*, can be used to perform the conversion.

► *Database characteristics*

Several items are available for determining database characteristics, such as its size and major and minor ODS version numbers. The following table lists the request buffer items that can be passed, and the information returned in the result buffer for each item type

Request buffer item	Result buffer contents
<i>allocation</i>	Number of database pages allocated
<i>base_level</i>	Database version (level) number: <ul style="list-style-type: none"> • 1 byte containing the number 1 • 1 byte containing the version number
<i>db_id</i>	<ul style="list-style-type: none"> • Database file name and site name: • 1 byte containing the number 2 for a local connection or 4 for a remote connection • 1 byte containing the length, <i>d</i>, of the database file name in bytes • A string of <i>d</i> bytes, containing the database file name • 1 byte containing the length, <i>l</i>, of the site name in bytes • A string of <i>l</i> bytes, containing the site name
<i>implementation</i>	Database implementation number: <ul style="list-style-type: none"> • 1 byte containing a 1 • 1 byte containing the implementation number • 1 byte containing a "class" number, either 1 or 12

TABLE 4.6 Database information items for database characteristics

Request buffer item	Result buffer contents
<i>no_reserve</i>	0 or 1 <ul style="list-style-type: none"> • 0 indicates space is reserved on each database page for holding backup versions of modified records [Default] • 1 indicates no space is reserved for such records
<i>ods_minor_version</i>	On-disk structure (ODS) minor version number; an increase in a minor version number indicates a non-structural change, one that still allows the database to be accessed by database engines with the same major version number but possibly different minor version numbers
<i>ods_version</i>	ODS major version number; databases with different major version numbers have different physical layouts A database engine can access only databases with a particular ODS major version number; trying to attach to a database with a different ODS number results in an error
<i>page_size</i>	Number of bytes per page of the attached database; use with <i>isc_info_allocation</i> to determine the size of the database
<i>version</i>	Version identification string of the database implementation: <ul style="list-style-type: none"> • 1 byte containing the number 1 • 1 byte specifying the length, <i>n</i>, of the following string • <i>n</i> bytes containing the version identification string

TABLE 4.6 Database information items for database characteristics (continued)

► *Environmental characteristics*

Several items are provided for determining environmental characteristics, such as the amount of memory currently in use, or the number of database cache buffers currently allocated. These items are described in the following table:

Request buffer item	Result buffer contents
<i>current_memory</i>	Amount of server memory (in bytes) currently in use
<i>forced_writes</i>	Number specifying the mode in which database writes are performed (0 for asynchronous, 1 for synchronous)
<i>max_memory</i>	Maximum amount of memory (in bytes) used at one time since the first process attached to the database
<i>num_buffers</i>	Number of memory buffers currently allocated
<i>sweep_interval</i>	Number of transactions that are committed between “sweeps” to remove database record versions that are no longer needed
<i>user_names</i>	Names of all the users currently attached to the database; for <i>each</i> such user, the result buffer contains an <i>isc_info_user_names</i> byte followed by a 1-byte length specifying the number of bytes in the user name, followed by the user name

TABLE 4.7 Database information items for environmental characteristics

Note Not all environmental information items are available on all platforms.

► *Performance statistics*

There are four items that request performance statistics for a database. These statistics accumulate for a database from the moment it is first attached by any process until the last remaining process detaches from the database.

For example, the value returned for *isc_info_reads* is the number of reads since the current database was first attached, that is, an *aggregate* of all reads done by all attached processes, rather than the number of reads done for the calling program since it attached to the database.

Table 4.8 summarizes the request performance statistics:

Request buffer item	Result buffer contents
<i>fetches</i>	Number of reads from the memory buffer cache
<i>marks</i>	Number of writes to the memory buffer cache
<i>reads</i>	Number of page reads
<i>writes</i>	Number of page writes

TABLE 4.8 Database information items for performance statistics

► *Database operation counts*

Several information items are provided for determining the number of various database operations performed by the currently attached calling program. These values are calculated on a per-table basis.

When any of these information items is requested, InterBase returns to the result buffer:

- 1 byte specifying the item type (for example, *isc_info_insert_count*).
- 2 bytes telling how many bytes compose the subsequent value pairs.
- A pair of values for each table in the database on which the requested type of operation has occurred since the database was last attached.

Each pair consists of:

- 2 bytes specifying the table ID.
- 4 bytes listing the number of operations (for example, inserts) done on that table.

Tip To determine an actual table name from a table ID, query the system table, RDB\$RELATION.

The following table describes the items which return count values for operations on the database:

Request buffer item	Result buffer contents
<i>backout_count</i>	Number of removals of a version of a record
<i>delete_count</i>	Number of database deletes since the database was last attached
<i>expunge_count</i>	Number of removals of a record and all of its ancestors, for records whose deletions have been committed
<i>insert_count</i>	Number of inserts into the database since the database was last attached
<i>purge_count</i>	Number of removals of old versions of fully mature records (records that are committed, so that older ancestor versions are no longer needed)
<i>read_idx_count</i>	Number of reads done via an index since the database was last attached
<i>read_seq_count</i>	Number of sequential sequential table scans (row reads) done on each table since the database was last attached
<i>update_count</i>	Number of database updates since the database was last attached

TABLE 4.9 Database information items for operation counts

***isc_database_info()* call example**

The following code requests the page size and the number of buffers for the currently attached database, then examines the result buffer:

```
char db_items[] = {
    isc_info_page_size, isc_info_num_buffers,
    isc_info_end};
char res_buffer[40], *p, item;
int length;
SLONG page_size = 0L, num_buffers = 0L;
ISC_STATUS status_vector[20];

isc_database_info(
    status_vector,
    &handle, /* Set in previous isc_attach_database() call. */
```

```

        sizeof(db_items),
        db_items,
        sizeof(res_buffer),
        res_buffer);
if (status_vector[0] == 1 && status_vector[1]) {
    /* An error occurred. */
    isc_print_status(status_vector);
    return(1);
};
/* Extract the values returned in the result buffer. */
for (p = res_buffer; *p != isc_info_end ; ) {
    item = *p++
    length = isc_vax_integer(p, 2);
    p += 2;
    switch (item){
        case isc_info_page_size:
            page_size = isc_vax_integer(p, length);
            break;
        case isc_info_num_buffers:
            num_buffers = isc_vax_integer(p, length);
            break;
        default:
            break;
    }
    p += length;
};

```

Disconnecting from databases

When an application is finished accessing a database, and any changes are committed or rolled back, the application should disconnect from the database, release system resources allocated for the attachment, and set the database handle to zero with a call to *isc_detach_database()*.

isc_detach_database() requires two arguments: a pointer to the error status vector, and a pointer to the handle of the database from which to detach. For example, the following statement detaches from the database pointed to by the database handle, *db1*:

```
isc_detach_database(status_vector, &db1);
```

Each database to detach requires a separate call to *isc_detach_database()*.

Deleting a database

To remove a database from the system if it is no longer needed, use *isc_drop_database()*. This function permanently wipes out a database, erasing its data, metadata, and all of its supporting files, such as secondary files, shadow files, and write-ahead log files.

A database can only be deleted if it is previously attached with a call to *isc_attach_database()*. The call to *isc_attach_database()* establishes a database handle for the database. That handle must be passed in the call to *isc_drop_database()*.

For example, the following code deletes the database pointed to by the database handle, *db1*:

```
#include <ibase.h>
. . .
isc_db_handle db1;
char *str = "employee.gdb";
ISC_STATUS status_vector[20];
. . .
/* Set database handle to zero before attaching to a database. */
db1 = 0L;
/* Attach to the database. */
isc_attach_database(status_vector, strlen(str), str, &db1, 0, NULL);
if (status_vector[0] == 1 && status_vector[1])
{
    error_exit();
}
isc_drop_database(status_vector, &db1);
if (status_vector[0] == 1 && status_vector[1])
{
    error_exit();
}
. . .
```

Working with Transactions

This chapter describes how to set up a *transaction parameter buffer* (TPB) that contains parameters, how to set up and initialize transaction handles, and how to use the API functions that control transactions. It also explains how to retrieve a transaction ID.

All data definition and data manipulation in an application takes place in the context of one or more *transactions*, one or more statements that work together to complete a specific set of actions that must be treated as an atomic unit of work.

The following table summarizes the API functions most commonly used when working with transactions. Functions are listed in the order they typically appear in an application.

Function	Purpose
<i>isc_start_transaction()</i>	Starts a new transaction against one or more databases.; use a previously declared and populated TPB
<i>isc_commit_retaining()</i>	Commits a transaction's changes, and preserves the transaction context for further transaction processing
<i>isc_commit_transaction()</i>	Commits a transaction's changes, and ends the transaction
<i>isc_rollback_transaction()</i>	Rolls back a transaction's changes, and ends the transaction

TABLE 5.1 API transaction functions

In addition to these functions, the following table lists less frequently used API transaction functions in the order they typically appear when used:

Function	Purpose
<i>isc_start_multiple()</i>	Starts a new transaction against one or more databases; used instead of <i>isc_start_transaction()</i> for programming languages such as FORTRAN, that do not support variable numbers of arguments to functions
<i>isc_prepare_transaction()</i>	Performs the first phase of a two-phase commit, prior to calling <i>isc_commit_transaction()</i> ; used only when it is absolutely necessary to override InterBase's automatic two-phase commit
<i>isc_prepare_transaction2()</i>	Performs the first phase of a two-phase commit, prior to calling <i>isc_commit_transaction()</i> ; used only when absolutely necessary to override InterBase's automatic two-phase commit

TABLE 5.2 Additional API transaction functions

Starting transactions

Starting transactions is a three-step process:

1. Creating and initializing a transaction handle for each simultaneous transaction to be started.
2. Optionally creating and populating a TPB for each transaction.

3. Calling `isc_start_transaction()` for each transaction to start.

These steps are described in the following sections of this chapter.

Note Programmers writing applications that do not permit function calls to pass a variable number of parameters must use `isc_start_multiple()` instead of `isc_start_transaction()`.

Creating transaction handles

Every transaction that is used in an application must be associated with its own *transaction handle*, a pointer to an address that is used by all API transaction functions. The `ibase.h` header file contains the following C typedef declaration for transaction handles:

```
typedef void ISC_FAR *isc_tr_handle;
```

To use this typedef for declaring transaction handles in an application, include `ibase.h` in each source file module:

```
#include <ibase.h>
```

► *Declaring transaction handles*

To establish transaction handles for use, declare a variable of type `isc_tr_handle` for each simultaneously active transaction. The following code declares two handles:

```
#include <ibase.h>
. . .
isc_tr_handle tr1;
isc_tr_handle tr2;
```

Once a transaction is committed or rolled back, its handle can be assigned to a different transaction in a subsequent call to `isc_start_transaction()`. If an application uses several transactions, but only starts a subset of transactions at the same time, it is only necessary to declare as many handles as there will be simultaneously active transactions. For example, if an application starts a total of three transactions, but only runs two of them at the same time, only two transaction handles need be declared.

► *Initializing transaction handles*

Before a transaction handle can be used to start a new transaction, it must be set to zero. The following code illustrates how two transaction handles are set to zero:

```
#include <ibase.h>
. . .
isc_tr_handle tr1;
```

```
isc_tr_handle tr2;
. . .
/* Set transaction handles to zero before starting a transaction. */
tr1 = 0L;
tr2 = 0L;
```

Once a transaction handle is initialized to zero, it can be used in a call to `isc_start_transaction()` to establish a new transaction. If a nonzero transaction handle is passed to `isc_start_transaction()`, the startup fails and an error code is returned. For more information about starting a new transaction with `isc_start_transaction()`, see **“Calling `isc_start_transaction()`” on page 70**.

Creating a transaction parameter buffer

The *transaction parameter buffer* (TPB) is an optional, application-defined byte vector, passed as an argument to `isc_start_transaction()`, that sets up a transaction’s *attributes*, its operating characteristics, such as whether the transaction has read and write access to tables, or read-only access, and whether or not other simultaneously active transactions can share table access with the transaction. Each transaction may have its own TPB, or transactions that share operating characteristics can use the same TPB.

Note If a TPB is not created for a transaction, a NULL pointer must be passed to `isc_start_transaction()` in its place. A default set of attributes is automatically assigned to such transactions. For more information about the default TPB, see **“Using the default TPB” on page 69**.

A TPB is declared in a C program as a char array of one-byte elements. Each element is a parameter that describes a single transaction attribute. A typical declaration is as follows:

```
static char isc_tpb[] = {isc_tpb_version3,
    isc_tpb_write,
    isc_tpb_read_committed,
    isc_tpb_no_rec_version,
    isc_tpb_wait};
```

This example makes use of parameter constants defined in the InterBase header file, **ibase.h**. The first element in every TPB must be the *isc_tpb_version3* constant. The following table lists available TPB constants, describes their purposes, and indicates which constants are assigned as a default set of attributes when a NULL TPB pointer is passed to *isc_start_transaction()*:

Parameter	Description
<i>isc_tpb_version3</i>	InterBase version 3 transaction
<i>isc_tpb_consistency</i>	Table-locking transaction model
<i>isc_tpb_concurrency</i>	High throughput, high concurrency transaction with acceptable consistency; use of this parameter takes full advantage of the InterBase multi-generational transaction model [Default]
<i>isc_tpb_shared</i>	Concurrent, shared access of a specified table among all transactions; use in conjunction with <i>isc_tpb_lock_read</i> and <i>isc_tpb_lock_write</i> to establish the lock option [Default]
<i>isc_tpb_protected</i>	Concurrent, restricted access of a specified table; use in conjunction with <i>isc_tpb_lock_read</i> and <i>isc_tpb_lock_write</i> to establish the lock option
<i>isc_tpb_wait</i>	Lock resolution specifies that the transaction is to wait until locked resources are released before retrying an operation [Default]
<i>isc_tpb_nowait</i>	Lock resolution specifies that the transaction is not to wait for locks to be released, but instead, a lock conflict error should be returned immediately
<i>isc_tpb_read</i>	Read-only access mode that allows a transaction only to select data from tables
<i>isc_tpb_write</i>	Read-write access mode of that allows a transaction to select, insert, update, and delete table data [Default]
<i>isc_tpb_lock_read</i>	Read-only access of a specified table. Use in conjunction with <i>isc_tpb_shared</i> , <i>isc_tpb_protected</i> , and <i>isc_tpb_exclusive</i> to establish the lock option.
<i>isc_tpb_lock_write</i>	Read-write access of a specified table. Use in conjunction with <i>isc_tpb_shared</i> , <i>isc_tpb_protected</i> , and <i>isc_tpb_exclusive</i> to establish the lock option [Default]

TABLE 5.3 TPB constants

Parameter	Description
<i>isc_tpb_read_committed</i>	High throughput, high concurrency transaction that can read changes committed by other concurrent transactions. Use of this parameter takes full advantage of the InterBase multi-generational transaction model.
<i>isc_tpb_rec_version</i>	Enables an <i>isc_tpb_read_committed</i> transaction to read the most recently committed version of a record even if other, uncommitted versions are pending.
<i>isc_tpb_no_rec_version</i>	Enables an <i>isc_tpb_read_committed</i> transaction to read only the latest committed version of a record. If an uncommitted version of a record is pending and <i>isc_tpb_wait</i> is also specified, then the transaction waits for the pending record to be committed or rolled back before proceeding. Otherwise, a lock conflict error is reported at once.

TABLE 5.3 TPB constants (continued)

TPB parameters specify the following classes of information:

- *Transaction version number* is used internally by the InterBase engine. It is always be the first attribute specified in the TPB, and must always be set to *isc_tpb_version3*.
- *Access mode* describes the actions that can be performed by the functions associated with the transaction. Valid access modes are:

isc_tpb_read

isc_tpb_write

- *Isolation level* describes the view of the database given a transaction as it relates to actions performed by other simultaneously occurring transactions. Valid isolation levels are:

isc_tpb_concurrency

isc_tpb_consistency

isc_tpb_read_committed, isc_tpb_rec_version

isc_tpb_read_committed, isc_tpb_no_rec_version

- *Lock resolution* describes how a transaction should react if a lock conflict occurs. Valid lock resolutions are:

isc_tpb_wait

isc_tpb_nowait

- *Table reservation* optionally describes an access method and lock resolution for a specified table that the transaction accesses. When table reservation is used, tables are reserved for the specified access when the transaction is started, rather than when the transaction actually accesses the table. Valid reservations are:

isc_tpb_shared, isc_tpb_lock_write

isc_tpb_shared, isc_tpb_lock_read

isc_tpb_protected, isc_tpb_lock_write

isc_tpb_protected, isc_tpb_lock_read

TPB parameters are described in detail in the following sections.

▶ *Specifying the transaction version number*

The first parameter in a TPB must always specify the version number for transaction processing. It must always be set to *isc_tpb_version3*. The following TPB declaration illustrates the correct use and position of this parameter:

```
static char isc_tpb[] = {isc_tpb_version3, ...};
```

▶ *Specifying access mode*

The access mode parameter describes the actions a transaction can perform against a table. The default access mode, *isc_tpb_write*, enables a transaction to read data from a table and write data to it. A second access mode, *isc_tpb_read*, restricts table access to read only. For example, the following TPB declaration specifies a read-only transaction:

```
static char isc_tpb[] = {isc_tpb_version3, isc_tpb_read};
```

A TPB should specify only one access mode parameter. If more than one is specified, later declarations override earlier ones.

If a TPB is declared that omits the access mode parameter, InterBase interprets transaction access as read and write.

▶ *Specifying isolation level*

The isolation level parameter specifies the view of the database permitted a transaction as it relates to actions performed by other simultaneously occurring transactions.

ISC_TPB_CONCURRENCY

By default, after a transaction starts it cannot access committed changes to a table made by other simultaneous transactions, even though it shares access to the table with them. Such a transaction has an isolation level of *isc_tpb_concurrency*, meaning it can have concurrent access to tables also accessed simultaneously by other transactions. The following declaration creates a TPB specifying an isolation level of *isc_tpb_concurrency*:

```
static char isc_tpb[] = {isc_tpb_version3,
    isc_tpb_write,
    isc_tpb_concurrency};
```

ISC_TPB_READ_COMMITTED

A second isolation level, *isc_tpb_read_committed*, offers all the advantages of the *isc_tpb_concurrency* isolation level and additionally enables a transaction to access changes committed by other simultaneous transactions. Two other parameters, *isc_tpb_rec_version*, and *isc_tpb_no_rec_version*, should be used with the *isc_tpb_read_committed* parameter. They offer refined control over the committed changes a transaction is permitted to access:

- *isc_tpb_no_rec_version*, the default refinement, specifies that a transaction can only read the latest version of a row. If a change to a row is pending, but not yet committed, the row cannot be read.
- *isc_tpb_rec_version* specifies that a transaction can read the latest committed version of a row, even if a more recent uncommitted version is pending.

The following declaration creates a TPB with a read committed isolation level, and specifies that the transaction can read the latest committed version of a row:

```
static char isc_tpb[] = {isc_tpb_version3,
    isc_tpb_write,
    isc_tpb_read_committed,
    isc_tpb_rec_version};
```

ISC_TPB_CONSISTENCY

InterBase also supports a restrictive isolation level. *isc_tpb_consistency* prevents a transaction from accessing tables if they are written to by other transactions; it also prevents other transactions from writing to a table once this transaction writes to it. This isolation level is designed to guarantee that if a transaction writes to a table before other simultaneous read and write transactions, then only it can change a table's data. Because it essentially restricts (and often prevents) shared access to tables, *isc_tpb_consistency* should be used with care.

A TPB should only specify one isolation mode parameter (and one refinement parameter, if isolation mode is *isc_tpb_read_committed*). If more than one is specified, later declarations override earlier ones.

If a TPB is declared that omits the isolation mode parameter, InterBase interprets it as *isc_tpb_concurrency*.

ISOLATION LEVEL INTERACTIONS

To determine the possibility for lock conflicts between two transactions accessing the same database, each transaction’s isolation level and access mode must be considered. The following table summarizes possible combinations:

		isc_tpb_concurrency, isc_tpb_read_committed		isc_tpb_consistency	
		isc_tpb_write	isc_tpb_read	isc_tpb_write	isc_tpb_read
concurrency, read_committed	<i>isc_tpb_write</i>	Some simultaneous updates may conflict	—	Conflicts	Conflicts
	<i>isc_tpb_read</i>	—	—	—	—
consistency	<i>isc_tpb_write</i>	Conflicts	—	Conflicts	Conflicts
	<i>isc_tpb_read</i>	Conflicts	—	Conflicts	—

TABLE 5.4 Isolation level interaction with read and write operations

As this table illustrates, *isc_tpb_concurrency* and *isc_tpb_read_committed* transactions offer the least chance for conflicts. For example, if *t1* is an *isc_tpb_concurrency* transaction with *isc_tpb_write* access, and *t2* is an *isc_tpb_read_committed* transaction with *isc_tpb_write* access, *t1* and *t2* only conflict when they attempt to update the same rows. If *t1* and *t2* have *isc_tpb_read* access, they never conflict with other transactions.

An *isc_tpb_consistency* transaction with *isc_tpb_write* access is guaranteed that if it gains access to a table that it alone can update a table, but it conflicts with all other simultaneous transactions except for *isc_tpb_concurrency* and *isc_tpb_read_committed* transactions running in *isc_tpb_read* mode. An *isc_tpb_consistency* transaction with *isc_tpb_read* access is compatible with any other read-only transaction, but conflicts with any transaction that attempts to insert, update, or delete data.

► *Specifying lock resolution*

The lock resolution parameter describes what happens if a transaction encounters an access conflict during a write operation (update and delete operations on existing rows). There are two possible choices for this parameter:

- *isc_tpb_wait*, the default, specifies that the transaction should wait until locked resources are released. Once the resources are released, the transaction retries its operation.
- *isc_tpb_nowait* specifies that the transaction should return a lock conflict error without waiting for locks to be released.

For example, the following declaration creates a TPB with write access, a concurrency isolation mode, and a lock resolution of *isc_tpb_nowait*:

```
static char isc_tpb[] = {isc_tpb_version3,
    isc_tpb_write,
    isc_tpb_concurrency,
    isc_tpb_nowait};
```

A TPB should only specify one lock resolution parameter. If more than one is specified, later declarations override earlier ones.

If a TPB is declared that omits the lock resolution parameter, InterBase interprets it as *isc_tpb_concurrency*.

► *Specifying table reservation*

Ordinarily, transactions gain specific access to tables only when they actually read from or write to them. Optional table reservation parameters can be passed in the TPB. Table reservation optionally describes an access method and lock resolution for a specified table that the transaction accesses. When table reservation is used, tables are reserved for the specified access when the transaction is started, rather than when the transaction actually accesses the table. Table reservation is only useful in an environment where simultaneous transactions share database access. It has three main purposes:

- Prevent possible deadlocks and update conflicts that can occur if locks are taken only when actually needed (the default behavior).
- Provide for *dependency locking*, the locking of tables that may be affected by triggers and integrity constraints. While explicit dependency locking is not required, it can assure that update conflicts do not occur because of indirect table conflicts.
- Change the level of shared access for one or more individual tables in a transaction. For example, an *isc_tpb_write* transaction with an isolation level of *isc_tpb_concurrency* may need exclusive update rights for a single table, and could use a reservation parameter to guarantee itself sole write access to the table.

Valid reservations are:

- *isc_tpb_shared*, *isc_tpb_lock_write*, which permits any transaction with an access mode of *isc_tpb_write* and isolation levels of *isc_tpb_concurrency* or *isc_tpb_read_committed*, to update, while other transactions with these isolation levels and an access mode of *isc_tpb_read* can read data.

- *isc_tpb_shared*, *isc_tpb_lock_read*, which permits any transaction to read data, and any transaction with an access mode of *isc_tpb_write* to update. This is the most liberal reservation mode.
- *isc_tpb_protected*, *isc_tpb_lock_write*, which prevents other transactions from updating. Other transactions with isolation levels of *isc_tpb_concurrency* or *isc_tpb_read_committed* can read data, but only this transaction can update.
- *isc_tpb_protected*, *isc_tpb_lock_read*, which prevents all transactions from updating, but permits all transactions to read data.

The name of the table to reserve must immediately follow the reservation parameters. For example, the following TPB declaration reserves a table, EMPLOYEE, for protected read access:

```
static char isc_tpb[] = {isc_tpb_version3,
    isc_tpb_write,
    isc_tpb_concurrency,
    isc_tpb_nowait,
    isc_tpb_protected, isc_tpb_lock_read, "EMPLOYEE"};
```

Several tables can be reserved at the same time. The following declaration illustrates how two tables are reserved, one for protected read, the other for protected write:

```
static char isc_tpb[] = {isc_tpb_version3,
    isc_tpb_write,
    isc_tpb_concurrency,
    isc_tpb_nowait,
    isc_tpb_protected, isc_tpb_lock_read, "COUNTRY",
    isc_tpb_protected, isc_tpb_lock_write, "EMPLOYEE"};
```

► *Using the default TPB*

Providing a TPB for a transaction is optional. If one is not provided, then a NULL pointer must be passed to *isc_start_transaction()* in place of a pointer to the TPB. In this case, InterBase treats a transaction as if the following TPB had been declared for it:

```
static char isc_tpb[] = {isc_tpb_version3,
    isc_tpb_write,
    isc_tpb_concurrency,
    isc_tpb_wait};
```

Calling *isc_start_transaction()*

Once transaction handles and TPBs are prepared, a transaction can be started by calling *isc_start_transaction()* using the following syntax:

```
ISC_STATUS isc_start_transaction(
    ISC_STATUS *status_vector,
    isc_tr_handle *trans_handle,
    short db_count,
    isc_db_handle *&db_handle,
    unsigned short tpb_length,
    char *tpb_ad);
```

For a transaction that runs against a single database, set *db_count* to 1. *db_handle* should be a database handle set with a previous call to *isc_attach_database()*. *tpb_length* is the size of the TPB passed in the next parameter, and *tpb_ad* is the address of the TPB. The following code illustrates a typical call to *isc_start_transaction()*:

```
#include <ibase.h>
. . .
ISC_STATUS status_vector[20];
isc_db_handle db1;
isc_tr_handle tr1;
static char isc_tpb[] = {isc_tpb_version3,
    isc_tpb_write,
    isc_tpb_concurrency,
    isc_tpb_wait};
. . .
/* Initialize database and transaction handles here. */
db1 = 0L;
tr1 = 0L;
. . .
/* Code for attaching to database here is omitted. */
isc_start_transaction(status_vector,
    &tr1,
    1,
    &db1,
    (unsigned short) sizeof(isc_tpb),
    isc_tpb);
```

A transaction can be opened against multiple databases. To do so, set the *db_count* parameter to the number of databases against which the transaction runs, then for each database, repeat the *db_handle*, *tpb_length*, and *tpb_ad* parameters as a group once for each database. For example, the following code fragment assumes that two databases are connected when the transaction is started:

```
isc_start_transaction(status_vector,
&tr1,
    2,
    &db1,
    (unsigned short) sizeof(isc_tpb),
    &tpb);
&db2,
    (unsigned short) sizeof(isc_tpb),
    &tpb);
```

For the complete syntax of *isc_start_transaction()*, see [page 386](#).

Calling *isc_start_multiple()*

An alternate method for starting a transaction against multiple databases is to use *isc_start_multiple()*. Using *isc_start_multiple()* is not recommended unless you:

- Are using a language that does not support a variable number of arguments in a function call.
- Do not know how many databases you want to attach to when coding the start of a transaction.

C programmers should seldom need to use this function.

isc_start_multiple() passes information about each target database to InterBase through an array of transaction existence blocks (TEBs). There must be one TEB for each database against which a transaction runs. A TEB is a structure you must declare in your applications as follows:

```
typedef struct {
    long *db_ptr;
    long tpb_len;
    char *tpb_ptr;
} ISC_TEB;
```

db_ptr is a pointer to a previously declared, initialized, and populated database handle. *tpb_len* is the size, in bytes, of the transaction parameter buffer (TPB) to use for the database, and *tpb_ptr* is a pointer to the TPB itself. For information about declaring, initializing, and populating a database handle, see **“Creating database handles” on page 42**. For more information about creating and populating a TPB, see **“Creating a transaction parameter buffer” on page 62**.

To use a TEB structure in an application, declare an array variable of type `ISC_TEB`. The number of array dimensions should correspond to the number of databases that the transaction runs against. For example, the following declaration creates an array of two TEBs, capable of supporting two databases:

```
ISC_TEB teb_array[2];
```

Once an array of TEBs is declared, and corresponding TBP are created and populated for each database, values may be assigned to the appropriate fields in the TEBs. For example, the following code illustrates how two TEBs are filled:

```
. . .
ISC_STATUS status_vector[20];
isc_db_handle db1, db2;
isc_tr_handle trans;
ISC_TEB teb_array[2];
. . .
db1 = db2 = 0L;
trans = 0L;
/* Code assumes that two TPBs, isc_tpb1, and isc_tpb2, are created
here. */
/* Code assumes databases are attached here. */
/* assign values to TEB array */
teb_array[0].db_ptr = &db1;
teb_array[0].tpb_len = sizeof(isc_tpb1);
teb_array[0].tpb_ptr = isc_tpb1;
teb_array[1].db_ptr = &db2;
teb_array[1].tpb_len = sizeof(isc_tpb2);
teb_array[1].tpb_ptr = isc_tpb2;
. . .
```

After the TEBs are loaded with values, *isc_start_multiple()* can be called using the following syntax:

```
ISC_STATUS isc_start_multiple(
    ISC_STATUS *status_vector,
    isc_tr_handle *trans_handle,
    short db_handle_count,
    void *teb_vector_address);
```

For example, the following statements starts a two-database transaction:

```
. . .
ISC_STATUS status_vector[20];
isc_db_handle db1, db2;
isc_tr_handle trans;
ISC_TEB teb_array[2];
. . .
db1 = db2 = 0L;
trans = 0L;
/* Code assumes that two TPBs, isc_tpb1, and isc_tpb2, are created
here. */
/* Code assumes databases are attached here. */
/* assign values to TEB array */
teb_array[0].db_ptr = &db1;
teb_array[0].tpb_len = sizeof(isc_tpb1);
teb_array[0].tpb_ptr = isc_tpb1;
teb_array[1].db_ptr = &db2;
teb_array[1].tpb_len = sizeof(isc_tpb2);
teb_array[1].tpb_ptr = isc_tpb2;
/* Start the transaction */
isc_start_multiple(status_vector, &trans, 2, teb_array);
. . .
```

Ending transactions

When a transaction's tasks are complete, or an error prevents a transaction from completing, the transaction must be ended to set the database to a consistent state. There are two API functions that end transactions:

- *isc_commit_transaction()* makes a transaction's changes permanent in the database. For transactions that span databases, this function performs an automatic, two-phase commit to ensure that all changes are made successfully.

- *isc_rollback_transaction()* undoes a transaction's changes, returning the database to its previous state, before the transaction started. This function is typically used when one or more errors occur that prevent a transaction from completing successfully.

Both *isc_commit_transaction()* and *isc_rollback_transaction()* close the record streams associated with the transaction, reinitialize the transaction name to zero, and release system resources allocated for the transaction. Freed system resources are available for subsequent use by any application or program.

isc_rollback_transaction() is frequently used inside error-handling routines to clean up transactions when errors occur. It can also be used to roll back a partially completed transaction prior to retrying it, and it can be used to restore a database to its prior state if a program encounters an unrecoverable error.

The API offers three additional functions for controlling transactions:

- *isc_commit_retaining()* commits a transaction but retains the current transaction's context—the system resources and cursor states used in the transaction—without requiring the overhead of ending a transaction, starting a new one, and reestablishing cursor states. In a busy, multi-user environment, maintaining transaction context for each user speeds up processing and uses fewer system resources than closing a transaction and opening a new one.
- *isc_prepare_transaction()* and *isc_prepare_transaction2()* enable an application to perform the first phase of an automatic, two-phase commit in its own time, then issue a call to *isc_commit_transaction()* to complete the commit.

IMPORTANT If the program ends before a transaction ends, a transaction is automatically rolled back, but databases are not closed. If a program ends without closing the database, data loss or corruption is possible. Therefore, open databases should always be closed by issuing an explicit call to *isc_detach_database()*.

For more information about detaching from a database, see **Chapter 4, “Working with Databases.”**

Using *isc_commit_transaction()*

Use *isc_commit_transaction()* to write transaction changes permanently to a database. *isc_commit_transaction()* closes the record streams associated with the transaction, resets the transaction name to zero, and frees system resources assigned to the transaction for other uses. The complete syntax for *isc_commit_transaction()* is:

```
ISC_STATUS isc_commit_transaction(
    ISC_STATUS *status_vector,
    isc_tr_handle *trans_handle);
```

For example, the following call commits a transaction:

```
isc_commit_transaction(status_vector, &trans);
```

where *status_vector* is a pointer to a previously declared error status vector, and *trans* is a pointer to a previously declared and initialized transaction handle.

Tip Even transactions started with an access mode of *isc_tpb_read* should be ended with a call to *isc_commit_transaction()* rather than *isc_rollback_transaction()*. The database is not changed, but the overhead required to start subsequent transactions is greatly reduced.

► Using *isc_commit_retaining()*

To write transaction changes to the database without establishing a new *transaction context*—the names, system resources, and current state of cursors used in a transaction—use *isc_commit_retaining()* instead of *isc_commit_transaction()*. In a busy, multi-user environment, maintaining the transaction context for each user speeds up processing and uses fewer system resources than closing and starting a new transaction for each action. The complete syntax for *isc_commit_retaining()* is:

```
ISC_STATUS isc_commit_retaining(
    ISC_STATUS *status_vector,
    isc_tr_handle *trans_handle);
```

isc_commit_retaining() writes all pending changes to the database, ends the current transaction *without* closing its record stream and cursors and without freeing its system resources, then starts a new transaction and assigns the existing record streams and system resources to the new transaction.

For example, the following call commits a specified transaction, preserving the current cursor status and system resources:

```
isc_commit_retaining(status_vector, &trans);
```

where *status_vector* is a pointer to a previously declared error status vector, and *trans* is a pointer to a previously declared and initialized transaction handle.

A call to *isc_rollback_transaction()* issued after *isc_commit_retaining()* only rolls back updates and writes occurring *after* the call to *isc_commit_retaining()*.

► *Using isc_prepare_transaction()*

When a transaction is committed against multiple databases using *isc_commit_transaction()*, InterBase automatically performs a two-phase commit. During the first phase of the commit, the InterBase engine polls all database participants to make sure they are still available, writes a message describing the transaction to the RDB\$TRANSACTION_DESCRIPTION field of the RDB\$TRANSACTION system table, then puts the transaction into a limbo state. It is during the second phase that transaction changes are actually committed to the database.

Some applications may have their own, additional requirements to make of the two-phase commit. These applications can call *isc_prepare_transaction()* to execute the first phase of the two-phase commit, then perform their own, additional tasks before completing the commit with a call to *isc_commit_transaction()*.

The syntax for *isc_prepare_transaction()* is:

```
ISC_STATUS isc_prepare_transaction(
    ISC_STATUS *status_vector,
    isc_tr_handle *trans_handle);
```

For example, the following code fragment illustrates how an application might call *isc_prepare_transaction()*, then its own routines, before completing a commit with *isc_commit_transaction()*:

```
ISC_STATUS status_vector[20];
isc_db_handle db1;
isc_tr_handle trans;
. . .
/* Initialize handles. */
db1 = 0L;
trans = 0L;
. . .
/* Code assumes a database is attached here, */
/* and a transaction started. */
. . .
/* Perform first phase of two-phase commit. */
isc_prepare_transaction(status_vector, &trans);
/* Application does its own processing here. */
my_app_function();
/* Now complete the two-phase commit. */
isc_commit_transaction(status_vector, &trans);
```

IMPORTANT It is generally a dangerous practice to delay the second phase of the commit after completing the first, because delays increase the chance that network or server problems can occur between phases.

Using *isc_prepare_transaction2()*

Like *isc_prepare_transaction()*, *isc_prepare_transaction2()* performs the first phase of a two-phase commit, except that *isc_prepare_transaction2()* enables an application to supply its own transaction description for insertion into the RDB\$TRANSACTION_DESCRIPTION field of the RDB\$TRANSACTION system table.

IMPORTANT Do not use this call without first examining and understanding the information InterBase stores in RDB\$TRANSACTION_DESCRIPTION during an automatic, two-phase commit. Storage of improper or incomplete information can prevent database recovery if the two-phase commit fails.

See [page 363](#) for the complete syntax of *isc_prepare_transaction2()*.

Using *isc_rollback_transaction()*

Use *isc_rollback_transaction()* to restore the database to its condition prior to the start of the transaction. *isc_rollback_transaction()* also closes the record streams associated with the transaction, resets the transaction name to zero, and frees system resources assigned to the transaction for other uses. *isc_rollback_transaction()* typically appears in error-handling routines. The syntax for *isc_rollback_transaction()* is:

```
ISC_STATUS isc_rollback_transaction(
    ISC_STATUS *status_vector,
    isc_tr_handle *trans_handle);
```

For example, the following call rolls back a transaction:

```
isc_rollback_transaction(status_vector, &trans);
```

where *status_vector* is a pointer to a previously declared error status vector, and *trans* is a pointer to a previously declared and initialized transaction handle.

Working with Dynamic SQL

This chapter describes how to use API dynamic SQL (DSQL) functions to handle dynamically created SQL statements for data definition and manipulation. Using low-level API calls enables client applications to build SQL statements or solicit them from end users at runtime, providing end users with a familiar database interface. It also provides applications developers low-level access to InterBase features, such as multiple databases, not normally available at a higher level with embedded DSQL statements. For example, the InterBase **isql** utility is a DSQL application built on low-level API calls.

All API DSQL function names begin with “isc_dsql” to make it easier to distinguish them from other API calls.

Overview of the DSQL programming process

Building and executing DSQL applications with the API involve the following general steps:

- Embedding DSQL API functions in an application.
- Using host-language facilities, such as datatypes and macros, to provide input and output areas for passing statements and parameters at runtime.

- Programming methods that use these statements and facilities to process SQL statements at runtime.

These steps are described in detail throughout this chapter.

DSQL API limitations

Although DSQL offers many advantages, it also has the following limitations:

- Dynamic transaction processing is not permitted; all named transactions must be declared at compile time.
- Dynamic access to Blob and array data is not supported; Blob and array data can be accessed, but only through standard, statically processed SQL statements, or through low-level API calls.
- Database creation is restricted to CREATE DATABASE statements executed within the context of EXECUTE IMMEDIATE.

For more information about database access in DSQL, see **“Accessing databases” on page 80**. For more information about handling transactions in DSQL applications, see **“Handling transactions” on page 81**. For more information about working with Blob data in DSQL, see **“Processing Blob data” on page 83**. For more information about handling array data in DSQL, see **“Processing array data” on page 83**. For more information about dynamic creation of databases, see **“Creating a database” on page 82**.

Accessing databases

The InterBase API permits applications to attach to multiple databases simultaneously using database handles. Database handles must be declared and initialized when an application is compiled. Separate database handles should be supplied and initialized for each database accessed simultaneously. For example, the following code creates a single handle, *db1*, and initializes it to zero:

```
#include <ibase.h>
isc_db_handle db1;
. . .
db1 = 0L;
```

Once declared and initialized, a database handle can be assigned dynamically to a database at runtime as follows:

```
#include <ibase.h>
```

```

. . .
char dbname[129];
ISC_STATUS status_vector[20];
. . .
prompt_user("Name of database to open: ");
gets(dbname);
isc_attach_database(status_vector, 0, dbname, &dbl, NULL, NULL);

```

A database handle can be used to attach to different databases as long as a previously attached database is first detached with *isc_detach_database()*, which automatically sets database handles to NULL. The following statements detach from a database, set the database handle to zero, and attach to a new database:

```

isc_detach_database(status_vector, &dbl);
isc_attach_database(status_vector, 0, "employee.gdb", &dbl, NULL,
NULL);

```

For more information about API function calls for databases, see **Chapter 4, “Working with Databases.”**

Handling transactions

InterBase requires that all transaction handles be declared when an application is compiled. Once fixed at compile time, transaction handles cannot be changed at runtime, nor can new handles be declared dynamically at runtime. Most API functions that process SQL statements at runtime, such as *isc_dsql_describe()*, *isc_dsql_describe_bind()*, *isc_dsql_execute()*, *isc_dsql_execute2()*, *isc_dsql_execute_immediate()*, *isc_dsql_exec_immed2()*, and *isc_dsql_prepare()*, support the inclusion of a transaction handle parameter. The SQL statements processed by these functions cannot pass transaction handles even if the SQL syntax for the statement permits the use of a TRANSACTION clause.

Before a transaction handle can be used, it must be declared and initialized to zero. The following code declares, initializes, and uses a transaction handle in an API call that allocates and prepares an SQL statement for execution:

```

#include <ibase.h>
. . .
isc_tr_handle trans; /* Declare a transaction handle. */
isc_stmt_handle stmt; /* Declare a statement handle. */
char *sql_stmt = "SELECT * FROM EMPLOYEE";
isc_db_handle dbl;
ISC_STATUS status_vector[20];
. . .

```

```

trans = 0L; /* Initialize the transaction handle to zero. */
stmt = NULL; /* Set handle to NULL before allocation. */
/* This code assumes both that a database attachment is made, */
/* and a transaction is started here. */
. . .
/* Allocate the SQL statement handle. */
isc_dsqli_allocate_statement(status_vector, &dbl, &stmt);
/* Prepare the statement for execution. */
isc_dsqli_prepare(status_vector, &trans, &stmt, 0, sql_stmt, 1, NULL);

```

Note The SQL SET TRANSACTION statement cannot be prepared with *isc_dsqli_prepare()*, but it can be processed with *isc_dsqli_execute_immediate()* if:

1. Previous transactions are first committed or rolled back.
2. The transaction handle is set to NULL.

For more information about using SQL statements, see the *Embedded SQL Guide*. For more information about SQL statement syntax, see the *Language Reference*.

Creating a database

To create a new database in an API application:

1. Detach from any currently attached databases with *isc_detach_database()*. Detaching from a database automatically sets its database handle to NULL.
2. Build the CREATE DATABASE statement to process.
3. Execute the statement with *isc_dsqli_execute_immediate()* or *isc_dsqli_exec_immed2()*.

For example, the following statements disconnect from any currently attached databases, and create a new database. Any existing database handles are set to NULL, so that they can be used to connect to the new database in future DSQL statements.

```

char *str = "CREATE DATABASE \"new_emp.gdb\"";
. . .
isc_detach_database(status_vector, &dbl);
isc_dsqli_execute_immediate(status_vector, &dbl, &trans, 0, str, 1,
    NULL);

```

Processing Blob data

Blob processing is not directly supported using DSQL, nor are Blob cursors supported. Applications that process SQL statements can use API calls to handle Blob processing. For more information about processing Blob data, see **Chapter 7, “Working with Blob Data.”**

Processing array data

Array processing is not directly supported using DSQL. DSQL applications can use API calls to process array data. For more information about array calls, see **Chapter 8, “Working with Array Data.”**

Writing an API application to process SQL statements

Writing an API application that processes SQL statements enables a developer to code directly to InterBase at a low level, while presenting end users a familiar SQL interface. API SQL applications are especially useful when any of the following are not known until runtime:

- The text of the SQL statement
- The number of host variables
- The datatypes of host variables
- References to database objects

Writing an API DSQL application is more complex than programming embedded SQL applications with regular SQL because for most DSQL operations, the application needs explicitly to allocate and process an extended SQL descriptor area (*XSQLDA*) data structure to pass data to and from the database.

To use the API to process a DSQL statement, follow these basic steps:

1. Determine if API calls can process the SQL statement.
2. Represent the SQL statement as a character string in the application.
3. If necessary, allocate one or more *XSQLDAs* for input parameters and return values.
4. Use appropriate API programming methods to process the SQL statement.

Determining if API calls can process an SQL statement

Except as noted earlier in this chapter, DSQL functions can process most SQL statements. For example, DSQL can process data manipulation statements such as DELETE and INSERT, data definition statements such as ALTER TABLE and CREATE INDEX, and SELECT statements.

The following table lists SQL statements that cannot be processed by DSQL functions:

Statement	Statement
CLOSE	DECLARE CURSOR
DESCRIBE	EXECUTE
EXECUTE IMMEDIATE	FETCH
OPEN	PREPARE

TABLE 6.1 SQL statements that cannot be processed by the API

These statements are used to process DSQL requests or to handle SQL cursors, which must always be specified when an application is written. Attempting to use them with DSQL results in run-time errors.

Representing an SQL statement as a character string

Within a DSQL application, an SQL statement can come from different sources. It might come directly from a user who enters a statement at a prompt, as does `isql`. Or it might be generated by the application in response to user interaction. Whatever the source of the SQL statement, it must be represented as an *SQL statement string*, a character string that is passed to DSQL for processing.

SQL statement strings do not begin with the EXEC SQL prefix or end with a semicolon (;) as they do in typical embedded applications. For example, the following host-language variable declaration is a valid SQL statement string:

```
char *str = "DELETE FROM CUSTOMER WHERE CUST_NO = 256";
```

Note The semicolon that appears at the end of this char declaration is a C terminator, and not part of the SQL statement string.

Specifying parameters in SQL statement strings

SQL statement strings often include *value parameters*, expressions that evaluate to a single numeric or character value. Parameters can be used anywhere in statement strings where SQL expects a value that is not the name of a database object.

A value parameter in a statement string can be passed as a constant, or passed as a placeholder at runtime. For example, the following statement string passes 256 as a constant:

```
char *str = "DELETE FROM CUSTOMER WHERE CUST_NO = 256";
```

It is also possible to build strings at runtime from a combination of constants. This method is useful for statements where the variable is not a true constant, or it is a table or column name, and where the statement is executed only once in the application.

To pass a parameter as a placeholder, the value is passed as a question mark (?) embedded within the statement string:

```
char *str = "DELETE FROM CUSTOMER WHERE CUST_NO = ?";
```

When a DSQL function processes a statement containing a placeholder, it replaces the question mark with a value supplied in an extended SQL descriptor area (*XSQLDA*) previously declared and populated in the application. Use placeholders in statements that are prepared once, but executed many times with different parameter values.

Replaceable value parameters are often used to supply values in SQL SELECT statement WHERE clause comparisons and in the UPDATE statement SET clause.

Understanding the XSQLDA


All DSQL applications must declare one or more extended SQL descriptor areas (*XSQLDAs*). The *XSQLDA* structure definition can be found in the **ibase.h** header file in the InterBase **include** directory. Applications declare instances of the *XSQLDA* for use.

The *XSQLDA* is a host-language data structure that DSQL uses to transport data to or from a database when processing an SQL statement string. There are two types of *XSQLDAs*: *input* descriptors and *output* descriptors. Both input and output descriptors are implemented using the *XSQLDA* structure.

One field in the *XSQLDA*, *sqlvar*, is an *XSQLVAR* structure. The *sqlvar* is especially important, because one *XSQLVAR* must be defined for each input parameter or column returned. Like the *XSQLDA*, the *XSQLVAR* is a structure defined in **ibase.h** in the InterBase **include** directory.

Applications do not declare instances of the *XSQLVAR* ahead of time, but must, instead, dynamically allocate storage for the proper number of *XSQLVAR* structures required for each DSQL statement before it is executed, then deallocate it, as appropriate, after statement execution.

The following figure illustrates the relationship between the *XSQLDA* and the *XSQLVAR*:



An input *XSQLDA* consists of a single *XSQLDA* structure and one *XSQLVAR* structure for each input parameter. An output *XSQLDA* also consists of one *XSQLDA* structure and one *XSQLVAR* structure for each data item returned by the statement. An *XSQLDA* and its associated *XSQLVAR* structures are allocated as a single block of contiguous memory.

The *isc_dsql_prepare()*, *isc_dsql_describe()*, and *isc_dsql_describe_bind()* functions can be used to determine the proper number of *XSQLVAR* structures to allocate, and the *XSQLDA_LENGTH* macro can be used to allocate the proper amount of space. For more information about the *XSQLDA_LENGTH* macro, see **“Using the XSQLDA_LENGTH macro” on page 90.**

XSQLDA field descriptions

The following table describes the fields that comprise the *XSQLDA* structure:

Field definition	Description
<i>short version</i>	Indicates the version of the <i>XSQLDA</i> structure. Set by an application. The current version is defined in <i>ibase.h</i> as <i>SQLDA_VERSION1</i>
<i>char sqldaid[8]</i>	Reserved for future use
<i>ISC_LONG sqldabc</i>	Reserved for future use
<i>short sqln</i>	Indicates the number of elements in the <i>sqlvar</i> array; the application should set this field whenever it allocates storage for a descriptor
<i>short sqld</i>	Indicates the number of parameters for an input <i>XSQLDA</i> , or the number of select-list items for an output <i>XSQLDA</i> ; set by InterBase during an <i>isc_dsql_describe()</i> , <i>isc_dsql_describe_bind()</i> , or <i>isc_dsql_prepare()</i> For an input descriptor, an <i>sqld</i> of 0 indicates that the SQL statement has no parameters; for an output descriptor, an <i>sqld</i> of 0 indicates that the SQL statement is not a SELECT statement
<i>XSQLVAR sqlvar</i>	The array of <i>XSQLVAR</i> structures; the number of elements in the array is specified in the <i>sqln</i> field

TABLE 6.2 *XSQLDA* field descriptions

The following table describes the fields that comprise the *XSQLVAR* structure:

Field definition	Description
<i>short sqltype</i>	Indicates the SQL datatype of parameters or select-list items; set by InterBase during <i>isc_dsql_describe()</i> , <i>isc_dsql_describe_bind()</i> , or <i>isc_dsql_prepare()</i>
<i>short sqlscale</i>	Provides scale, specified as a negative number, for exact numeric datatypes (DECIMAL, NUMERIC); set by InterBase during <i>isc_dsql_describe()</i> , <i>isc_dsql_describe_bind()</i> , or <i>isc_dsql_prepare()</i>
<i>short sqlsubtype</i>	Specifies the subtype for Blob data; set by InterBase during <i>isc_dsql_describe()</i> , <i>isc_dsql_describe_bind()</i> , or <i>isc_dsql_prepare()</i>
<i>short sqllen</i>	Indicates the maximum size, in bytes, of data in the <i>sqldata</i> field; set by InterBase during <i>isc_dsql_describe()</i> , <i>isc_dsql_describe_bind()</i> , or <i>isc_dsql_prepare()</i>
<i>char *sqldata</i>	For input descriptors, specifies either the address of a select-list item or a parameter; set by the application For output descriptors, contains a value for a select-list item; set by InterBase
<i>short *sqlind</i>	On input, specifies the address of an indicator variable; set by an application; on output, specifies the address of column indicator value for a select-list item following a FETCH A value of 0 indicates that the column is not NULL; a value of -1 indicates the column is NULL; set by InterBase
<i>short sqlname_length</i>	Specifies the length, in bytes, of the data in field, <i>sqlname</i> ; set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>
<i>char sqlname[32]</i>	Contains the name of the column. Not NULL (\0) terminated; set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>
<i>short rename_length</i>	Specifies the length, in bytes, of the data in field, <i>rename</i> ; set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>

TABLE 6.3 *XSQLVAR* field descriptions

Field definition	Description
<i>char relname[32]</i>	Contains the name of the table; not NULL (\0) terminated, set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>
<i>short ownname_length</i>	Specifies the length, in bytes, of the data in field, <i>ownname</i> ; set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>
<i>char ownname[32]</i>	Contains the name of the table owner; not NULL (\0) terminated, set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>
<i>short aliasname_length</i>	Specifies the length, in bytes, of the data in field, <i>aliasname</i> ; set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>
<i>char aliasname[32]</i>	Contains the alias name of the column. If no alias exists, contains the column name; not NULL (\0) terminated, set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>

TABLE 6.3 XSQLVAR field descriptions (continued)

Input descriptors

Input descriptors are used to process SQL statement strings that contain parameters. Before an application can execute a statement with parameters, it must supply values for them. The application indicates the number of parameters passed in the *XSQLDA sqlld* field, then describes each parameter in a separate *XSQLVAR* structure. For example, the following statement string contains two parameters, so an application must set *sqlld* to 2, and describe each parameter:

```
char *str = "UPDATE DEPARTMENT SET BUDGET = ? WHERE LOCATION = ?";
```

When the statement is executed, the first *XSQLVAR* supplies information about the BUDGET value, and the second *XSQLVAR* supplies the LOCATION value.

For more information about using input descriptors, see **“DSQL programming methods” on page 96**.

Output descriptors

Output descriptors return values from an executed query to an application. The *sqlld* field of the *XSQLDA* indicates how many values were returned. Each value is stored in a separate *XSQLVAR* structure. The *XSQLDA sqlvar* field points to the first of these *XSQLVAR* structures. The following statement string requires an output descriptor:

```
char *str = "SELECT * FROM CUSTOMER WHERE CUST_NO > 100";
```

For information about retrieving information from an output descriptor, see **“DSQL programming methods” on page 96**.

Using the `XSQLDA_LENGTH` macro

The `ibase.h` header file defines a macro, `XSQLDA_LENGTH`, to calculate the number of bytes that must be allocated for an input or output `XSQLDA`. `XSQLDA_LENGTH` is defined as follows:

```
#define XSQLDA_LENGTH (n) (sizeof (XSQLDA) + (n - 1) * sizeof(XSQLVAR))
```

n is the number of parameters in a statement string, or the number of select-list items returned from a query. For example, the following C statement uses the `XSQLDA_LENGTH` macro to specify how much memory to allocate for an `XSQLDA` with 5 parameters or return items:

```
XSQLDA *my_xsqlda;
. . .
my_xsqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(5));
. . .
```

For more information about using the `XSQLDA_LENGTH` macro, see **“DSQL programming methods” on page 96**.

SQL datatype macro constants

InterBase defines a set of macro constants to represent SQL datatypes and NULL status information in an `XSQLVAR`. An application should use these macro constants to specify the datatype of parameters and to determine the datatypes of select-list items in an SQL statement. The following table lists each SQL datatype, its corresponding macro constant expression, C datatype or InterBase typedef, and whether or not the `sqlind` field is used to indicate a parameter or variable that contains NULL or unknown data:

SQL datatype	Macro expression	C datatype or typedef	<code>sqlind</code> used?
Array	<code>SQL_ARRAY</code>	<code>ISC_QUAD</code>	No
Array	<code>SQL_ARRAY + 1</code>	<code>ISC_QUAD</code>	Yes
Blob	<code>SQL_BLOB</code>	<code>ISC_QUAD</code>	No

TABLE 6.4 SQL datatypes, macro expressions, and C datatypes

SQL datatype	Macro expression	C datatype or typedef	<i>sqliind</i> used?
BLOB	SQL_BLOB + 1	ISC_QUAD	Yes
CHAR	SQL_TEXT	char[]	No
CHAR	SQL_TEXT + 1	char[]	Yes
DATE	SQL_DATE	ISC_DATE	No
DATE	SQL_DATE + 1	ISC_DATE	Yes
DECIMAL	SQL_SHORT, SQL_LONG, SQL_DOUBLE, or SQL_INT64	int, long, double, or ISC_INT64	No
DECIMAL	SQL_SHORT + 1, SQL_LONG + 1, SQL_DOUBLE + 1, or SQL_INT64 + 1	int, long, double, or ISC_INT64	Yes
DOUBLE PRECISION	SQL_DOUBLE	double	No
DOUBLE PRECISION	SQL_DOUBLE + 1	double	Yes
INTEGER	SQL_LONG	long	No
INTEGER	SQL_LONG + 1	ISC_LONG	Yes
FLOAT	SQL_FLOAT	float	No
FLOAT	SQL_FLOAT + 1	float	Yes
NUMERIC	SQL_SHORT, SQL_LONG, SQL_DOUBLE, or SQL_INT64	int, long, double, or ISC_INT64	No
NUMERIC	SQL_SHORT + 1, SQL_LONG + 1, SQL_DOUBLE + 1, or SQL_INT64 + 1	int, long, double, or ISC_INT64	Yes
SMALLINT	SQL_SHORT	short	No
SMALLINT	SQL_SHORT + 1	short	Yes
TIME	SQL_TIME	ISC_TIME	No
TIME	SQL_TIME + 1	ISC_TIME	Yes
TIMESTAMP	SQL_TIMESTAMP	ISC_TIMESTAMP	No

TABLE 6.4 SQL datatypes, macro expressions, and C datatypes (*continued*)

SQL datatype	Macro expression	C datatype or typedef	<i>sqlind</i> used?
TIMESTAMP	SQL_TIMESTAMP + 1	ISC_TIMESTAMP	Yes
VARCHAR	SQL_VARYING	First 2 bytes: short containing the length of the character string; remaining bytes: char[]	No
VARCHAR	SQL_VARYING + 1	First 2 bytes: short containing the length of the character string; remaining bytes: char[]	Yes

TABLE 6.4 SQL datatypes, macro expressions, and C datatypes (*continued*)

Note DECIMAL and NUMERIC datatypes are stored internally as SMALLINT, INTEGER, DOUBLE PRECISION, or 64-bit integer datatypes. To specify the correct macro expression to provide for a DECIMAL or NUMERIC column, use **isql** to examine the column definition in the table to see how InterBase is storing column data, then choose a corresponding macro expression.

The datatype information for a parameter or select-list item is contained in the *sqltype* field of the *XSQLVAR* structure. The value contained in *sqltype* provides two pieces of information:

- The datatype of the parameter or select-list item.
- Whether *sqlind* is used to indicate NULL values. If *sqlind* is used, its value specifies whether the parameter or select-list item is NULL (–1), or not NULL (0).

For example, if *sqltype* equals SQL_TEXT, the parameter or select-list item is a CHAR that does not use *sqlind* to check for a NULL value (because, in theory, NULL values are not allowed for it). If *sqltype* equals SQL_TEXT + 1, then *sqlind* can be checked to see if the parameter or select-list item is NULL.

Tip The C language expression, *sqltype & 1*, provides a useful test of whether a parameter or select-list item can contain a NULL. The expression evaluates to 0 if the parameter or select-list item cannot contain a NULL, and 1 if the parameter or select-list item can contain a NULL. The following code fragment demonstrates how to use the expression:

```
if (sqltype & 1 == 0)
{
    /* parameter or select-list item that CANNOT contain a NULL */
}
else
{
    /* parameter or select-list item CAN contain a NULL */
}
```

By default, both `isc_dsql_prepare()` and `isc_dsql_describe()` return a macro expression of type + 1, so `sqlind` should always be examined for NULL values with these statements.

Handling varying string datatypes

VARCHAR, CHARACTER VARYING, and NCHAR VARYING datatypes require careful handling in DSQL. The first two bytes of these datatypes contain string length information, while the remainder of the data contains the actual bytes of string data to process.

To avoid having to write code to extract and process variable-length strings in an application, it is possible to force these datatypes to fixed length using SQL macro expressions. For more information about forcing variable-length data to fixed length for processing, see **“Coercing datatypes” on page 94**.

Applications can, instead, detect and process variable-length data directly. To do so, they must extract the first two bytes from the string to determine the byte-length of the string itself, then read the string, byte-by-byte, into a null-terminated buffer.

Handling NUMERIC and DECIMAL datatypes

DECIMAL and NUMERIC datatypes are stored internally as SMALLINT, INTEGER, DOUBLE PRECISION, or 64-bit integer datatypes, depending on the precision and scale defined for a column definition that uses these types. To determine how a DECIMAL or NUMERIC value is actually stored in the database, use `isql` to examine the column definition in the table. If NUMERIC is reported, then data is actually being stored as DOUBLE PRECISION.

When a DECIMAL or NUMERIC value is stored as a SMALLINT, INTEGER, or 64-bit integer, the value is stored as a whole number. During retrieval in DSQL, the `sqlscale` field of the `XSQLVAR` is set to a negative number that indicates the factor of 10 by which the whole number (returned in `sqldata`), must be divided in order to produce the correct NUMERIC or DECIMAL value with its fractional part. If `sqlscale` is -1, then the number must be divided by 10, if it is -2, then the number must be divided by 100, -3 by 1000, and so forth.

Coercing datatypes

Sometimes when processing DSQL input parameters and select-list items, it is desirable or necessary to translate one datatype to another. This process is referred to as *datatype coercion*. For example, datatype coercion is often used when parameters or select-list items are of type VARCHAR. The first two bytes of VARCHAR data contain string length information, while the remainder of the data is the string to process. By coercing the data from SQL_VARYING to SQL_TEXT, data processing can be simplified.

Coercion can only be from one compatible datatype to another. For example, SQL_VARYING to SQL_TEXT, or SQL_SHORT to SQL_LONG.

► Coercing character datatypes

To coerce SQL_VARYING datatypes to SQL_TEXT datatypes, change the *sqltype* field in the parameter's or select-list item's *XSQLVAR* structure to the desired SQL macro datatype constant. For example, the following statement assumes that *var* is a pointer to an *XSQLVAR* structure, and that it contains an SQL_VARYING datatype to convert to SQL_TEXT:

```
var->sqltype = SQL_TEXT;
```

After coercing a character datatype, provide proper storage space for it. The *XSQLVAR* field, *sqlen*, contains information about the size of the uncoerced data.

Set the *XSQLVAR sqldata* field to the address of the data.

► Coercing numeric datatypes

To coerce one numeric datatype to another, change the *sqltype* field in the parameter's or select-list item's *XSQLVAR* structure to the desired SQL macro datatype constant. For example, the following statement assumes that *var* is a pointer to an *XSQLVAR* structure, and that it contains an SQL_SHORT datatype to convert to SQL_LONG:

```
var->sqltype = SQL_LONG;
```

IMPORTANT Do not coerce a larger datatype to a smaller one. Data can be lost in such a translation.

► Setting a NULL indicator

If a parameter or select-list item contains a NULL value, the *sqlind* field should be used to indicate its NULL status. Appropriate storage space must be allocated for *sqlind* before values can be stored there.

Before insertion, set *sqlind* to -1 to indicate that NULL values are legal. Otherwise, set *sqlind* to 0.

After selection, an *sqlind* of -1 indicates a field contains a NULL value. Other values indicate a field contains non-NULL data.

Aligning numerical data

Ordinarily, when a variable with a numeric datatype is created, the compiler will ensure that the variable is stored at a properly aligned address, but when numeric data is stored in a dynamically allocated buffer space, such as can be pointed to by the *XSQLDA* and *XSQLVAR* structures, the programmer must take precautions to ensure that the storage space is properly aligned.

Certain platforms, in particular those with RISC processors, require that numerical data in dynamically allocated storage structures be aligned properly in memory. Alignment is dependent both on datatype and platform.

For example, a short integer on a Sun SPARCstation must be located at an address divisible by 2, while a long on the same platform must be located at an address divisible by 4. In most cases, a data item is properly aligned if the address of its starting byte is divisible by the correct alignment number. Consult specific system and compiler documentation for alignment requirements.

A useful rule of thumb is that the size of a datatype is always a valid alignment number for the datatype. For a given type T, if size of (T) equals *n*, then addresses divisible by *n* are correctly aligned for T. The following macro expression can be used to align data:

```
#define ALIGN(ptr, n) ((ptr + n - 1) & ~(n - 1))
```

where *ptr* is a pointer to char.

The following code illustrates how the ALIGN macro might be used:

```
char *buffer_pointer, *next_aligned;
next_aligned = ALIGN(buffer_pointer, sizeof(T));
```

DSQL programming methods

There are four possible DSQL programming methods for handling an SQL statement string. The best method for processing a string depends on the type of SQL statement in the string, and whether or not it contains placeholders for parameters. The following decision table explains how to determine the appropriate processing method for a given string:

Is it a query?	Does it have placeholders?	Processing method to use:
No	No	Method 1
No	Yes	Method 2
Yes	No	Method 3
Yes	Yes	Method 4

TABLE 6.5 SQL statement strings and recommended processing methods

Method 1: Non-query statements without parameters

There are two ways to process an SQL statement string containing a non-query statement without placeholder parameters:

- Use *isc_dsql_execute_immediate()* to prepare and execute the string a single time.
- Use *isc_dsql_allocate_statement()* to allocate a statement string for the statement to execute, *isc_dsql_prepare()* to parse the statement for execution and assign it a name, then use *isc_dsql_execute()* to carry out the statement's actions as many times as required in an application.

► Using *isc_dsql_execute_immediate()*

1. To execute a statement string a single time, use *isc_dsql_execute_immediate()*:
2. Elicit a statement string from the user or create one that contains the SQL statement to be processed. For example, the following statement creates an SQL statement string:

```
char *str = "UPDATE DEPARTMENT SET BUDGET = BUDGET * 1.05";
```

3. Parse and execute the statement string using *isc_dsql_execute_immediate()*:

```
isc_dsql_execute_immediate(status_vector, &dbl, &trans,
    0, str, 1, NULL);
```

Note `isc_dsql_execute_immediate()` also accepts string literals. For example,

```
isc_dsql_execute_immediate(status_vector, &dbl, &trans, 0,
    "UPDATE DEPARTMENT SET BUDGET = BUDGET * 1.05", 1, NULL);
```

For the complete syntax of `isc_dsql_execute_immediate()` and an explanation of its parameters, see **Chapter 13, “API Function Reference.”**

► *Using `isc_dsql_prepare()` and `isc_dsql_execute()`*

To execute a statement string several times, use `isc_dsql_allocate_statement()`, `isc_dsql_prepare()`, and `isc_dsql_execute()`:

1. Elicit a statement string from the user or create one that contains the SQL statement to be processed. For example, the following statement creates an SQL statement string:

```
char *str = "UPDATE DEPARTMENT SET BUDGET = BUDGET * 1.05";
```

2. Declare and initialize an SQL statement handle, then allocate it with `isc_dsql_allocate_statement()`:

```
isc_stmt_handle stmt; /* Declare a statement handle. */
stmt = NULL; /* Set handle to NULL before allocation. */
. . .
isc_dsql_allocate_statement(status_vector, &dbl, &stmt);
```

3. Parse the statement string with `isc_dsql_prepare()`. This sets the statement handle (*stmt*) to refer to the parsed format. The statement handle is used in subsequent calls to `isc_dsql_execute()`:

```
isc_dsql_prepare(status_vector, &trans, &stmt, 0, str, 1, NULL);
```

Note `isc_dsql_prepare()` also accepts string literals. For example,

```
isc_dsql_prepare(status_vector, &trans, &stmt, 0,
    "UPDATE DEPARTMENT SET BUDGET = BUDGET * 1.05", 1, NULL);
```

4. Execute the named statement string using `isc_dsql_execute()`. For example, the following statement executes a statement string named *stmt*:

```
isc_dsql_execute(status_vector, &trans, &stmt, 1, NULL);
```

Once a statement string is prepared, it can be executed as many times as required in an application.

Method 2: Non-query statements with parameters

There are two steps to processing an SQL statement string containing a non-query statement with placeholder parameters:

1. Create an input *XSQLDA* to process a statement string's parameters.
2. Prepare and execute the statement string with its parameters.

► *Creating the input XSQLDA*

Placeholder parameters are replaced with actual data before a prepared SQL statement string is executed. Because those parameters are unknown when the statement string is created, an input *XSQLDA* must be created to supply parameter values at execute time. To prepare the *XSQLDA*, follow these steps:

1. Declare a variable to hold the *XSQLDA* needed to process parameters. For example, the following declaration creates an *XSQLDA* called *in_sqlda*:

```
XSQLDA *in_sqlda;
```

2. Optionally declare a variable for accessing the *XSQLVAR* structure of the *XSQLDA*:

```
XSQLVAR *var;
```

Declaring a pointer to the *XSQLVAR* structure is not necessary, but can simplify referencing the structure in subsequent statements.

3. Allocate memory for the *XSQLDA* using the *XSQLDA_LENGTH* macro. The following statement allocates storage for *in_sqlda*:

```
in_sqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(10));
```

In this statement space for 10 *XSQLVAR* structures is allocated, allowing the *XSQLDA* to accommodate up to 10 parameters.

4. Set the *version* field of the *XSQLDA* to *SQLDA_VERSION1*, and set the *sqln* field to indicate the number of *XSQLVAR* structures allocated:

```
in_sqlda->version = SQLDA_VERSION1;
in_sqlda->sqln = 10;
```

► *Preparing and executing a statement string with parameters*

After an *XSQLDA* is created for holding a statement string's parameters, the statement string can be created and prepared. Local variables corresponding to the placeholder parameters in the string must be assigned to their corresponding *sqldata* fields in the *XSQLVAR* structures.

To prepare and execute a non-query statement string with parameters, follow these steps:

1. Elicit a statement string from the user or create one that contains the SQL statement to be processed. For example, the following statement creates an SQL statement string with placeholder parameters:

```
char *str = "UPDATE DEPARTMENT SET BUDGET = ?, LOCATION = ?";
```

This statement string contains two parameters: a value to be assigned to the BUDGET column and a value to be assigned to the LOCATION column.

2. Declare and initialize an SQL statement handle, then allocate it with *isc_dsqli_allocate()*:

```
isc_stmt_handle stmt; /* Declare a statement handle. */
stmt = NULL; /* Set handle to NULL before allocation. */
. . .
isc_dsqli_allocate_statement(status_vector, &dbl, &stmt);
```

3. Parse the statement string with *isc_dsqli_prepare()*. This sets the statement handle (*stmt*) to refer to the parsed format. The statement handle is used in subsequent calls to *isc_dsqli_describe_bind()* and *isc_dsqli_execute()*:

```
isc_dsqli_prepare(status_vector, &trans, &stmt, 0, str, 1,
in_sqlda);
```

4. Use *isc_dsqli_describe_bind()* to fill the input *XSQLDA* with information about the parameters contained in the SQL statement:

```
isc_dsqli_describe_bind(status_vector, &stmt, 1, in_sqlda);
```

5. Compare the value of the *sqln* field of the *XSQLDA* to the value of the *sqld* field to make sure enough *XSQLVARs* are allocated to hold information about each parameter. *sqln* should be at least as large as *sqld*. If not, free the storage previously allocated to the input descriptor, reallocate storage to reflect the number of parameters specified by *sqld*, reset *sqln* and *version*, then execute *isc_dsqli_describe_bind()* again:

```
if (in_sqlda->sqld > in_sqlda->sqln)
{
    n = in_sqlda->sqld;
    free(in_sqlda);
    in_sqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(n));
    in_sqlda->sqln = n;
    in_sqlda->version = SQLDA_VERSION1;
    isc_dsqli_describe_bind(status_vector, &stmt, 1, in_sqlda);
}
```

6. Process each *XSQLVAR* parameter structure in the *XSQLDA*. Processing a parameter structure involves up to four steps:
 - Coerce a parameter's datatype (optional).
 - Allocate local storage for the data pointed to by the *sqldata* field of the *XSQLVAR*. This step is only required if space for local variables is not allocated until runtime. The following example illustrates dynamic allocation of local variable storage space.
 - Provide a value for the parameter consistent with its datatype (required).
 - Provide a NULL value indicator for the parameter.

The following code example illustrates these steps, looping through each *XSQLVAR* structure in the *in_sqlda XSQLDA*:

```
for (i=0, var = in_sqlda->sqlvar; i < in_sqlda->sqlda; i++, var++)
{
    /* Process each XSQLVAR parameter structure here.
    Var points to the parameter structure. */

    dtype = (var->sqltype & ~1) /* drop NULL flag for now */
    switch(dtype)
    {
        case SQL_VARYING: /* coerce to SQL_TEXT */
            var->sqltype = SQL_TEXT;
            /* allocate local variable storage */
            var->sqldata = (char *)malloc(sizeof(char)*var->sqlllen);
            . . .
            break;
        case SQL_TEXT:
            var->sqldata = (char *)malloc(sizeof(char)*var->sqlllen);
            /* provide a value for the parameter */
            . . .
            break;
        case SQL_LONG:
            var->sqldata = (char *)malloc(sizeof(long));
            /* provide a value for the parameter */
            *(long *) (var->sqldata) = 17;
            break;
        . . .
    } /* end of switch statement */
    if (sqltype & 1)
    {
        /* allocate variable to hold NULL status */
```

```

        var->sqlind = (short *)malloc(sizeof(short));
    }
} /* end of for loop */

```

For more information about datatype coercion and NULL indicators, see **“Coercing datatypes” on page 94**.

- Execute the named statement string with `isc_dsql_execute()`. For example, the following statement executes a statement string named `stmt`:

```
isc_dsql_execute(status_vector, &trans, &stmt, 1, in_sqlda);
```

► *Re-executing the statement string*

Once a non-query statement string with parameters is prepared, it can be executed as often as required in an application. Before each subsequent execution, the input `XSQLDA` can be supplied with new parameter and NULL indicator data.

To supply new parameter and NULL indicator data for a prepared statement, repeat step 6 of **“Preparing and executing a statement string with parameters” on page 98**.

Method 3: Query statements without parameters

There are three steps to processing an SQL query statement string without parameters:

- Prepare an output `XSQLDA` to process the select-list items returned when the query is executed.
- Prepare the statement string.
- Use a cursor to execute the statement and retrieve select-list items from the output `XSQLDA`.

► *Preparing the output XSQLDA*

Most queries return one or more rows of data, referred to as a *select-list*. Because the number and kind of items returned are unknown when a statement string is created, an output `XSQLDA` must be created to store select-list items that are returned at runtime. To prepare the `XSQLDA`, follow these steps:

- Declare a variable to hold the `XSQLDA` needed to store the column data for each row that will be fetched. For example, the following declaration creates an `XSQLDA` called `out_sqlda`:

```
XSQLDA *out_sqlda;
```

- Optionally declare a variable for accessing the *XSQLVAR* structure of the *XSQLDA*:

```
XSQLVAR *var;
```

Declaring a pointer to the *XSQLVAR* structure is not necessary, but can simplify referencing the structure in subsequent statements.

- Allocate memory for the *XSQLDA* using the *XSQLDA_LENGTH* macro. The following statement allocates storage for *out_sqlda*:

```
out_sqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(10));
```

Space for ten *XSQLVAR* structures is allocated in this statement, enabling the *XSQLDA* to accommodate up to ten select-list items.

- Set the *version* field of the *XSQLDA* to *SQLDA_VERSION1*, and set the *sqln* field of the *XSQLDA* to indicate the number of *XSQLVAR* structures allocated:

```
out_sqlda->version = SQLDA_VERSION1;
out_sqlda->sqln = 10;
```

► *Preparing a query statement string without parameters*

After an *XSQLDA* is created for holding the items returned by a query statement string, the statement string can be created, prepared, and described. When a statement string is executed, InterBase creates the select-list of selected rows.

To prepare a query statement string, follow these steps:

- Elicit a statement string from the user or create one that contains the SQL statement to be processed. For example, the following statement creates an SQL statement string that performs a query:

```
char *str = "SELECT * FROM CUSTOMER";
```

The statement appears to have only one select-list item (*). The asterisk is a wildcard symbol that stands for all of the columns in the table, so the actual number of items returned equals the number of columns in the table.

- Declare and initialize an SQL statement handle, then allocate it with *isc_dsqli_allocate()*:

```
isc_stmt_handle stmt; /* Declare a statement handle. */
stmt = NULL; /* Set handle to NULL before allocation. */
. . .
isc_dsqli_allocate_statement(status_vector, &dbl, &stmt);
```

3. Parse the statement string with `isc_dsql_prepare()`. This sets the statement handle (*stmt*) to refer to the parsed format. The statement handle is used in subsequent calls to statements such as `isc_dsql_describe()` and `isc_dsql_execute()`:

```
isc_dsql_prepare(status_vector, &trans, &stmt, 0, str, 1, NULL);
```

4. Use `isc_dsql_describe()` to fill the output *XSQLDA* with information about the select-list items returned by the statement:

```
isc_dsql_describe(status_vector, &trans, &stmt, out_sqlda);
```

5. Compare the *sqln* field of the *XSQLDA* to the *sqld* field to determine if the output descriptor can accommodate the number of select-list items specified in the statement. If not, free the storage previously allocated to the output descriptor, reallocate storage to reflect the number of select-list items specified by *sqld*, reset *sqln* and *version*, then execute `isc_dsql_describe()` again:

```
if (out_sqlda->sqld > out_sqlda->sqln)
{
    n = out_sqlda->sqld;
    free(out_sqlda);
    out_sqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(n));
    out_sqlda->sqln = n;
    out_sqlda->version = SQLDA_VERSION1;
    isc_dsql_describe(status_vector, &trans, 1, out_sqlda);
}
```

6. Set up an *XSQLVAR* structure for each item returned. Setting up an item structure involves the following steps:
 - Coercing an item's datatype (optional).
 - Allocating local storage for the data pointed to by the *sqldata* field of the *XSQLVAR*. This step is only required if space for local variables is not allocated until runtime. The following example illustrates dynamic allocation of local variable storage space.
 - Providing a NULL value indicator for the parameter.

The following code example illustrates these steps, looping through each *XSQLVAR* structure in the *out_sqlda XSQLDA*:

```
for (i=0, var = out_sqlda->sqlvar; i < out_sqlda->sqld; i++, var++)
{
    dtype = (var->sqltype & ~1) /* drop flag bit for now */
    switch(dtype)
    {
```

```

case SQL_VARYING:
    var->sqltype = SQL_TEXT;
    var->sqldata = (char *)malloc(sizeof(char)*var->sqlllen + 2);
    break;
case SQL_TEXT:
    var->sqldata = (char *)malloc(sizeof(char)*var->sqlllen);
    break;
case SQL_LONG:
    var->sqldata = (char *)malloc(sizeof(long));
    break;
    . . .
    /* process remaining types */
} /* end of switch statements */
if (sqltype & 1)
{
    /* allocate variable to hold NULL status */
    var->sqlind = (short *)malloc(sizeof(short));
}
} /* end of for loop */

```

For more information about datatype coercion and NULL indicators, see **“Coercing datatypes” on page 94**.

► *Executing a statement string within the context of a cursor*

To retrieve select-list items from a prepared statement string, the string can be executed within the context of a cursor. All cursor declarations in InterBase are fixed statements inserted into the application before it is compiled. DSQL application developers must anticipate the need for cursors when writing the application and declare them ahead of time.

A cursor is only needed to process positioned UPDATE and DELETE statements made against the rows retrieved by *isc_dsql_fetch()* for SELECT statements that specify an optional FOR UPDATE OF clause.

The following descriptions apply to the situations when a cursor is needed. For an example of executing a statement and fetching rows without using a cursor, see **“isc_dsql_fetch()” on page 330**.

A looping construct is used to fetch a single row at a time from the cursor and to process each select-list item (column) in that row before the next row is fetched.

To execute a statement string within the context of a cursor and retrieve rows of select-list items, follow these steps:

1. Execute the prepared statement with `isc_dsql_execute()`:

```
isc_dsql_execute(status_vector, &trans, &stmt, 1, NULL);
```

2. Declare and open a cursor for the statement string with `isc_dsql_set_cursor_name()`. For example, the following statement declares a cursor, `dyn_cursor`, for the SQL statement string, `stmt`:

```
isc_dsql_set_cursor_name(status_vector, &stmt,
    "dyn_cursor", NULL);
```

Opening the cursor causes the statement string to be executed, and an active set of rows to be retrieved.

3. Fetch one row at a time and process the select-list items (columns) it contains with `isc_dsql_fetch()`. For example, the following loops retrieve one row at a time from `dyn_cursor` and process each item in the retrieved row with an application-specific function called `process_column()`:

```
while ((fetch_stat =
    isc_dsql_fetch(status_vector, &stmt, 1, out_sqlda))
    == 0)
{
    for (i = 0; i < out_sqlda->sqlc; i++)
    {
        process_column(sqlda->sqlvar[i]);
    }
}
if (fetch_stat != 100L)
{
    /* isc_dsql_fetch returns 100 if no more rows remain to be
       retrieved */
    SQLCODE = isc_sqlcode(status_vector);
    isc_print_sqlerror(SQLCODE, status_vector);
    return(1);
}
```

The `process_column()` function mentioned in this example processes each returned select-list item. The following skeleton code illustrates how such a function can be set up:

```
void process_column(XSQLVAR *var)
{
    /* test for NULL value */
    if ((var->sqltype & 1) && (*(var->sqlind) = -1))
    {
        /* process the NULL value here */
    }
}
```

```

    }
    else
    {
        /* process the data instead */
    }
    . . .
}

```

4. When all the rows are fetched, close the cursor with

isc_dsql_free_statement():

```
isc_dsql_free_statement(status_vector, &stmt, DSQL_close);
```

► *Re-executing a query statement string without parameters*

Once a query statement string without parameters is prepared, it can be executed as often as required in an application by closing and reopening its cursor.

To reopen a cursor and process select-list items, repeat steps 2 through 4 of “**Executing a statement string within the context of a cursor**” on page 104.

Method 4: Query statements with parameters

There are four steps to processing an SQL query statement string with placeholder parameters:

1. Prepare an input *XSQLDA* to process a statement string’s parameters.
2. Prepare an output *XSQLDA* to process the select-list items returned when the query is executed.
3. Prepare the statement string and its parameters.
4. Use a cursor to execute the statement using input parameter values from an input *XSQLDA*, and to retrieve select-list items from the output *XSQLDA*.

► *Preparing the input XSQLDA*

Placeholder parameters are replaced with actual data before a prepared SQL statement string is executed. Because those parameters are unknown when the statement string is created, an input *XSQLDA* must be created to supply parameter values at runtime. To prepare the *XSQLDA*, follow these steps:

1. Declare a variable to hold the *XSQLDA* needed to process parameters. For example, the following declaration creates an *XSQLDA* called *in_sqlda*:

```
XSQLDA in_sqlda;
```

2. Optionally declare a variable for accessing the *XSQLVAR* structure of the *XSQLDA*:

```
XSQLVAR *var;
```

Declaring a pointer to the *XSQLVAR* structure is not necessary, but can simplify referencing the structure in subsequent statements.

3. Allocate memory for the *XSQLDA* using the *XSQLDA_LENGTH* macro. The following statement allocates storage for *in_sqlda*:

```
in_sqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(10));
```

In this statement, space for 10 *XSQLVAR* structures is allocated, allowing the *XSQLDA* to accommodate up to 10 input parameters. Once structures are allocated, assign values to the *sqldata* fields.

4. Set the *version* field of the *XSQLDA* to *SQLDA_VERSION1*, and set the *sqln* field of the *XSQLDA* to indicate the number of *XSQLVAR* structures allocated:

```
in_sqlda->version = SQLDA_VERSION1;
in_sqlda->sqln = 10;
```

► *Preparing the output XSQLDA*

Most queries return one or more rows of data, referred to as a *select-list*. Because the number and kind of items returned are unknown when a statement string is executed, an output *XSQLDA* must be created to store select-list items that are returned at runtime. To prepare the *XSQLDA*, follow these steps:

1. Declare a variable to hold the *XSQLDA* needed to process parameters. For example, the following declaration creates an *XSQLDA* called *out_sqlda*:

```
XSQLDA *out_sqlda;
```

2. Optionally declare a variable for accessing the *XSQLVAR* structure of the *XSQLDA*:

```
XSQLVAR *var;
```

Declaring a pointer to the *XSQLVAR* structure is not necessary, but can simplify referencing the structure in subsequent statements.

3. Allocate memory for the *XSQLDA* using the *XSQLDA_LENGTH* macro. The following statement allocates storage for *out_sqlda*:

```
out_sqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(10));
```

Space for ten *XSQLVAR* structures is allocated in this statement, enabling the *XSQLDA* to accommodate up to ten select-list items.

4. Set the *version* field of the *XSQLDA* to *SQLDA_VERSION1*, and set the *sqln* field of the *XSQLDA* to indicate the number of *XSQLVAR* structures allocated:

```
out_sqllda->version = SQLDA_VERSION1;
out_sqllda->sqln = 10;
```

► *Preparing a query statement string with parameters*

After an input and an output *XSQLDA* are created for holding a statement string's parameters, and the select-list items returned when the statement is executed, the statement string can be created and prepared. When a statement string is prepared, InterBase replaces the placeholder parameters in the string with information about the actual parameters used. The information about the parameters must be assigned to the input *XSQLDA* (and perhaps adjusted) before the statement can be executed. When the statement string is executed, InterBase stores select-list items in the output *XSQLDA*.

To prepare a query statement string with parameters, follow these steps:

1. Elicit a statement string from the user or create one that contains the SQL statement to be processed. For example, the following statement creates an SQL statement string with placeholder parameters:

```
char *str = "SELECT * FROM DEPARTMENT WHERE BUDGET = ?,
           LOCATION = ?";
```

This statement string contains two parameters: a value to be assigned to the BUDGET column and a value to be assigned to the LOCATION column.

2. Declare and initialize an SQL statement handle, then allocate it with *isc_dsql_allocate()*:

```
isc_stmt_handle stmt; /* Declare a statement handle. */
stmt = NULL; /* Set handle to NULL before allocation. */
. . .
isc_dsql_allocate_statement(status_vector, &dbl, &stmt);
```

3. Prepare the statement string with *isc_dsql_prepare()*. This sets the statement handle (*stmt*) to refer to the parsed format. The statement handle is used in subsequent calls to *isc_dsql_describe()*, *isc_dsql_describe_bind()*, and *isc_dsql_execute2()*:

```
isc_dsql_prepare(status_vector, &trans, &stmt, 0,
               str, 1, out_xsqllda);
```

4. Use *isc_dsql_describe_bind()* to fill the input *XSQLDA* with information about the parameters contained in the SQL statement:

```
isc_dsql_describe_bind(status_vector, &stmt, 1, in_xsqllda);
```

5. Compare the *sqln* field of the *XSQLDA* to the *sqld* field to determine if the input descriptor can accommodate the number of parameters contained in the statement. If not, free the storage previously allocated to the input descriptor, reallocate storage to reflect the number of parameters specified by *sqld*, reset *sqln* and *version*, then execute *isc_dsql_describe_bind()* again:

```

if (in_sqlda->sqld > in_sqlda->sqln)
{
    n = in_sqlda->sqld;
    free(in_sqlda);
    in_sqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(n));
    in_sqlda->sqln = n;
    in_sqlda->version = SQLDA_VERSION1;
    isc_dsql_describe_bind(status_vector, &stmt, 1, in_xsqlda);
}

```

6. Process each *XSQLVAR* parameter structure in the input *XSQLDA*. Processing a parameter structure involves up to four steps:
- Coercing a parameter's datatype (optional).
 - Allocating local storage for the data pointed to by the *sqldata* field of the *XSQLVAR*. This step is only required if space for local variables is not allocated until runtime. The following example illustrates dynamic allocation of local variable storage space.
 - Providing a value for the parameter consistent with its datatype (required).
 - Providing a NULL value indicator for the parameter.

These steps must be followed in the order presented. The following code example illustrates these steps, looping through each *XSQLVAR* structure in the *in_sqlda XSQLDA*:

```

for (i=0, var = in_sqlda->sqlvar; i < in_sqlda->sqld; i++, var++)
{
    /* Process each XSQLVAR parameter structure here.
    The parameter structure is pointed to by var.*/
    dtype = (var->sqltype & ~1) /* drop flag bit for now */
    switch(dtype)
    {
        case SQL_VARYING: /* coerce to SQL_TEXT */
            var->sqltype = SQL_TEXT;
            /* allocate proper storage */
            var->sqldata = (char *)malloc(sizeof(char)*var->sqllen);
            /* Provide a value for the parameter. See case SQL_LONG. */
            . . .
            break;

```

```

case SQL_TEXT:
    var->sqldata = (char *)malloc(sizeof(char)*var->sqlllen);
    /* Provide a value for the parameter. See case SQL_LONG. */
    . . .
    break;
case SQL_LONG:
    var->sqldata = (char *)malloc(sizeof(long));
    /* Provide a value for the parameter. */
    *(long *) (var->sqldata) = 17;
    break;
    . . .
} /* end of switch statement */
if (sqltype & 1)
{
    /* allocate variable to hold NULL status */
    var->sqlind = (short *)malloc(sizeof(short));
}
} /* end of for loop */

```

For more information about datatype coercion and NULL indicators, see **“Coercing datatypes” on page 94**.

7. Use `isc_dsql_describe()` to fill the output `XSQLDA` with information about the select-list items returned by the statement:

```
isc_dsql_describe(status_vector, &trans, &stmt, out_xsqlda);
```

8. Compare the `sqln` field of the `XSQLDA` to the `sqld` field to determine if the output descriptor can accommodate the number of select-list items specified in the statement. If not, free the storage previously allocated to the output descriptor, reallocate storage to reflect the number of select-list items specified by `sqld`, reset `sqln` and `version`, and execute DESCRIBE OUTPUT again:

```

if (out_sqlda->sqld > out_sqlda->sqln)
{
    n = out_sqlda->sqld;
    free(out_sqlda);
    out_sqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(n));
    out_sqlda->sqln = n;
    out_sqlda->version = SQLDA_VERSION1;
    isc_dsql_describe(status_vector, &trans, &stmt, out_xsqlda);
}

```

9. Set up an `XSQLVAR` structure for each item returned. Setting up an item structure involves the following steps:

- Coercing an item's datatype (optional).
- Allocating local storage for the data pointed to by the *sqldata* field of the *XSQLVAR*. This step is required only if space for local variables is not allocated until runtime. The following example illustrates dynamic allocation of local variable storage space.
- Providing a NULL value indicator for the parameter (optional).

The following code example illustrates these steps, looping through each *XSQLVAR* structure in the *out_sqlda XSQLDA*:

```
for (i=0, var = out_sqlda->sqlvar; i < out_sqlda->sqld; i++, var++)
{
    dtype = (var->sqltype & ~1) /* drop flag bit for now */
    switch(dtype)
    {
        case SQL_VARYING:
            var->sqltype = SQL_TEXT;
            break;
        case SQL_TEXT:
            var->sqldata = (char *)malloc(sizeof(char)*var->sqlllen);
            break;
        case SQL_LONG:
            var->sqldata = (char *)malloc(sizeof(long));
            break;
            /* process remaining types */
    } /* end of switch statements */
    if (sqltype & 1)
    {
        /* allocate variable to hold NULL status */
        var->sqlind = (short *)malloc(sizeof(short));
    }
} /* end of for loop */
```

For more information about datatype coercion and NULL indicators, see **“Coercing datatypes” on page 94**.

► *Executing a query statement string within the context of a cursor*

To retrieve select-list items from a statement string, the string must be executed within the context of a cursor. All cursor declarations in InterBase are fixed, embedded statements inserted into the application before it is compiled. DSQL application developers must anticipate the need for cursors when writing the application and declare them ahead of time.

A looping construct is used to fetch a single row at a time from the cursor and to process each select-list item (column) in that row before the next row is fetched.

To execute a statement string within the context of a cursor and retrieve rows of select-list items, follow these steps:

1. Execute the statement with *isc_dsqli_execute2()*:

```
isc_dsqli_execute2(status_vector, &trans, &stmt, 1,
                  in_xsqlda, out_xsqlda);
```

2. Declare and open a cursor for the statement string with *isc_dsqli_set_cursor_name()*. For example, the following statement declares a cursor, *dyn_cursor*, for the prepared SQL statement string, *stmt*:

```
isc_dsqli_set_cursor_name(status_vector, &stmt, "dyn_cursor", NULL);
```

Opening the cursor causes the statement string to be executed, and an active set of rows to be retrieved.

3. Fetch one row at a time with *isc_dsqli_fetch()* and process the select-list items (columns) it contains. For example, the following loops retrieve one row at a time from *dyn_cursor* and process each item in the retrieved row with an application-specific function called *process_column()*:

```
while ((fetch_stat =
       isc_dsqli_fetch(status_vector, &stmt, 1, out_sqlda)) == 0)
{
    for (i = 0; i < out_sqlda->sqld; i++)
    {
        process_column(sqlda->sqlvar[i]);
    }
}
if (fetch_stat != 100L)
{
    /* isc_dsqli_fetch returns 100 if no more rows remain to be
       retrieved */
    SQLCODE = isc_sqlcode(status_vector);
    isc_print_sqlerror(SQLCODE, status_vector);
    return(1);
}
```

4. When all the rows are fetched, close the cursor with *isc_dsqli_free_statement()*:

```
isc_dsqli_free_statement(status_vector, &stmt, DSQL_close);
```

► *Re-executing a query statement string with parameters*

Once a query statement string with parameters is prepared, it can be used as often as required in an application. Before each subsequent use, the input *XSQLDA* can be supplied with new parameter and NULL indicator data. The cursor must be closed and reopened before processing can occur.

- To provide new parameters to the input *XSQLDA*, follow steps 3 to 5 of **“Preparing a query statement string with parameters” on page 108.**
- To provide new information to the output *XSQLDA*, follow steps 6 to 8 of **“Preparing a query statement string with parameters” on page 108.**
- To reopen a cursor and process select-list items, repeat steps 2 to 4 of **“Executing a query statement string within the context of a cursor” on page 111.**

Determining an unknown statement type at runtime

An application can use *isc_dsql_sql_info()* to determine the statement type of an unknown prepared statement, for example, a statement entered by the user at runtime.

Requested information can include:

- Statement type.
- Number of input parameters required by the statement.
- Number of output values returned by the statement.
- Detailed information regarding each input parameter or output value, including its datatype, scale, and length.

To use *isc_dsql_sql_info()*, allocate an item-list buffer that describes the type of information requested, and allocate a result buffer, where the function can return the desired information. For example, to determine the statement type of an unknown, but prepared statement, you would allocate a one-element item-list buffer, and fill it with the macro constant, *isc_info_sql_stmt_type*, defined in **ibase.h**:

```
char type_item[];
type_item[] = {isc_info_sql_stmt_type};
```

Note Additional information item macros for requested items can be found in **ibase.h** under the comment, “SQL information items.”

The result buffer must be large enough to contain any data returned by the call. The proper size for this buffer depends on the information requested. If not enough space is allocated, then *isc_dsql_sql_info()* puts the predefined value, *isc_info_truncated*, in the last byte of the result buffer. Generally, when requesting statement type information, 8 bytes is a sufficient buffer size. Declaring a larger than necessary buffer is also safe. A request to identify a statement type returns the following information in the result buffer:

1. One byte containing *isc_info_sql_stmt_type*.
2. Two bytes containing a number, *n*, telling how many bytes compose the subsequent *value*.
3. One or two bytes specifying the statement type. The following table lists the statement types that can be returned:

Type	Numeric value
<i>isc_info_sql_stmt_select</i>	1
<i>isc_info_sql_stmt_insert</i>	2
<i>isc_info_sql_stmt_update</i>	3
<i>isc_info_sql_stmt_delete</i>	4
<i>isc_info_sql_stmt_ddl</i>	5
<i>isc_info_sql_stmt_get_segment</i>	6
<i>isc_info_sql_stmt_put_segment</i>	7
<i>isc_info_sql_stmt_exec_procedure</i>	8
<i>isc_info_sql_stmt_start_trans</i>	9
<i>isc_info_sql_stmt_commit</i>	10
<i>isc_info_sql_stmt_rollback</i>	11
<i>isc_info_sql_stmt_select_for_upd</i>	12

TABLE 6.6 Statement types

4. A final byte containing the value *isc_info_end (0)*.

The values placed in the result buffer are not aligned. Furthermore, all numbers are represented in a generic format, with the least significant byte first, and the most significant byte last. Signed numbers have the sign in the last byte. Convert the numbers to a datatype native to your system before interpreting them.

Note All information about a statement except its type can be more easily determined by calling functions other than *isc_dsqli_sql_info()*. For example, to determine the information to fill in an input *XSQLDA*, call *isc_dsqli_describe_bind()*. To fill in an output *XSQLDA*, call *isc_dsqli_prepare()* or *isc_dsqli_describe()*.

Working with Blob Data

This chapter describes InterBase's dynamically sizable datatype, called a Blob, and describes how to work with it using API functions. Depending on a particular application, you might need to read all or only part of the chapter.

For example, if you plan to request conversion of Blob data from one datatype to another, such as from one bitmapped graphic format to another or from the MIDI sound format to the Wave format, you need to read the entire chapter. To write a conversion routine, called a *filter*, see **“Filtering Blob data” on page 136**. For further information about working with Blob data and filters, see the *Embedded SQL Guide*.

Note Blob filters are not available on NetWare servers.

If you do not need to request conversion of Blob data, then you only need to read the parts of this chapter up to **“Filtering Blob data” on page 136**.

The following table alphabetically lists the API functions for working with Blob data. The functions will be described and demonstrated in the remainder of this chapter.

Function	Purpose
<i>isc_blob_default_desc()</i>	Loads a Blob descriptor with default information about a Blob, including its subtype, character set, and segment size
<i>isc_blob_gen_bpb()</i>	Generates a Blob parameter buffer (BPB) from source and target Blob descriptors to allow dynamic access to Blob subtype and character set information
<i>isc_blob_info()</i>	Returns information about an open Blob
<i>isc_blob_lookup_desc()</i>	Determines the subtype, character set, and segment size of a Blob, given a table name and Blob column name
<i>isc_blob_set_desc()</i>	Initializes a Blob descriptor from parameters passed to it
<i>isc_cancel_blob()</i>	Discards a Blob
<i>isc_close_blob()</i>	Closes an open Blob
<i>isc_create_blob2()</i>	Creates and opens a Blob for write access, and optionally specifies a filter to be used to translate the Blob from one subtype to another
<i>isc_get_segment()</i>	Retrieves data from a Blob column in a row returned by execution of a SELECT statement
<i>isc_open_blob2()</i>	Opens an existing Blob for retrieval, and optionally specifies a filter to be used to translate the Blob from one subtype to another
<i>isc_put_segment()</i>	Writes data into a Blob

TABLE 7.1 API Blob functions

What is a Blob?

A *Blob* is an object that cannot easily be stored in a database as one of the standard datatypes. You can use a Blob to store large amounts of data of various types, including:

- Bitmapped images
- Sounds
- Video segments
- Text

InterBase support of Blob data provides all the advantages of a database management system, including transaction control, maintenance, and access using standard API function calls. Blob data is stored in the database itself. Other systems only store pointers in the database to non-database files. InterBase stores the actual Blob data in the database, and establishes a unique identification handle in the appropriate table to point to the database location of the Blob. By maintaining the Blob data within the database, InterBase greatly improves access to and management of the data.

How are Blob data stored?

Blob is the InterBase datatype that can represent various objects, such as bitmapped images, sound, video, and text. Before you store these items in the database, you create or manage them as platform- or product-specific files or data structures, such as:

- TIFF, PICT, .BMP, .WMF, .GEM, TARGA or other bitmapped or vector-graphic files
- MIDI or .WAV sound files
- Audio Video Interleaved Format (.AVI) or QuickTime video files
- ASCII, .MIF, .DOC, .WPx or other text files
- CAD files

You must programmatically load these files from memory into the database, as you do any other data items or records you intend to store in InterBase. For more information about creating a Blob and storing data into it, see **“Writing data to a Blob” on page 126**.

Blob subtypes

Although you manage Blob data in ways similar to other datatypes, InterBase provides more flexible data typing rules for Blob data. Because there are many native datatypes that you can define as Blob data, InterBase treats them generically and allows you to define your own datatype, known as a *subtype*. InterBase also provides two predefined subtypes: 0, an unstructured subtype generally applied to binary data or data of an indeterminate type, and 1, applied to plain text.

User-defined subtypes must always be represented as negative integers between -1 and -32,678.

A Blob column's subtype is specified when the Blob column is defined.

The application is responsible for ensuring that data stored in a Blob column agrees with its subtype; InterBase does not check the type or format of Blob data.

Blob database storage

Rather than storing Blob data directly in the Blob field of a table record, InterBase stores a *Blob ID* there. A Blob ID is a unique numeric value that references Blob data. The Blob data is stored elsewhere in the database, in a series of *Blob segments*, units of Blob data read and written in chunks. Blob segments can be of varying length. The length of an individual segment is specified when it is written.

Segments are handy when working with data that is too large for one application memory buffer. But it is not necessary to use multiple segments; you can put all your Blob data in a single segment.

When an application creates a Blob, it must write data to it a segment at a time. When an application reads a Blob, it reads a segment at a time. For more information about writing segments, see **“Writing data to a Blob” on page 126**. For more information about reading segments, see **“Reading data from a Blob” on page 121**.

Blob data operations

InterBase supports the following operations on Blob data:

- Reading from a Blob
- Writing to a Blob, which involves the following operations:
 1. Inserting a new row that includes Blob data.
 2. Replacing the data referenced by a Blob column of a row.
 3. Updating the data referenced by a Blob column of a row.
- Deleting a Blob

The following sections describe how to perform these operations. These examples do not include the use of filters to convert data from one subtype to another as it is read or written. For information about using filters, see **“Writing an application that requests filtering” on page 142**.

Dynamic SQL (DSQL) API functions and the *XSQLDA* data structure are needed to execute SELECT, INSERT, and UPDATE statements required to select, insert, or update relevant Blob data. The following sections include descriptions of the DSQL programming methods required to execute the sample statements provided. For more information about DSQL programming, see **Chapter 6, “Working with Dynamic SQL.”**

Reading data from a Blob

There are six steps required for reading data from an existing Blob:

1. Create a `SELECT` statement query that specifies selection of the Blob column (and any other columns desired) in the rows of interest.
2. Prepare an output `XSQLDA` structure to hold the column data for each row that is fetched.
3. Prepare the `SELECT` statement for execution.
4. Execute the statement.
5. Fetch the selected rows one by one.
6. Read and processing the Blob data from each row.

► *Creating the SELECT statement*

Elicit a statement string from the user or create one that consists of the SQL query that will select rows containing the Blob data of interest. For example, the following creates an SQL query statement string that selects three columns from various rows in the `PROJECT` table:

```
char *str =
"SELECT PROJ_NAME, PROJ_DESC, PRODUCT FROM PROJECT WHERE \
PRODUCT IN ('software', 'hardware', 'other') ORDER BY PROJ_NAME";
```

► *Preparing the output XSQLDA*

Most queries return one or more rows of data, referred to as a *select-list*. An output `XSQLDA` must be created to store the column data for each row that is fetched. For a Blob column, the column data is an internal Blob identifier (Blob ID) that is needed to access the actual data. To prepare the `XSQLDA`, follow these steps:

1. Declare a variable to hold the `XSQLDA`. For example, the following declaration creates an `XSQLDA` called `out_sqlda`:

```
XSQLDA *out_sqlda;
```

2. Allocate memory for the `XSQLDA` using the `XSQLDA_LENGTH` macro. The `XSQLDA` must contain one `XSQLVAR` substructure for each column to be fetched. The following statement allocates storage for an output `XSQLDA` (`out_sqlda`) with three `XSQLVAR` substructures:

```
out_sqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(3));
```

3. Set the `version` field of the `XSQLDA` to `SQLDA_VERSION1`, and set the `sqln` field of the `XSQLDA` to indicate the number of `XSQLVAR` substructures allocated:

```

out_sqlda->version = SQLDA_VERSION1;
out_sqlda->sqln = 3;

```

► *Preparing the SELECT statement for execution*

After an *XSQLDA* is created for holding the column data for each selected row, the query statement string can be prepared for execution. Follow these steps:

1. Declare and initialize an SQL statement handle, then allocate it with *isc_dsql_allocate_statement()*:

```

isc_stmt_handle stmt; /* Declare a statement handle. */
stmt = NULL; /* Set handle to NULL before allocation. */
isc_dsql_allocate_statement(status_vector, &db_handle, &stmt);

```

2. Ready the statement string for execution with *isc_dsql_prepare()*. This checks the string (*str*) for syntax errors, parses it into a format that can be efficiently executed, and sets the statement handle (*stmt*) to refer to this parsed format. The statement handle is used in a later call to *isc_dsql_execute()*.

If *isc_dsql_prepare()* is passed a pointer to the output *XSQLDA*, as in the following example, it will fill in most fields of the *XSQLDA* and all its *XSQLVAR* substructures with information such as the datatype, length, and name of the corresponding columns in the statement.

A sample call to *isc_dsql_prepare()* is:

```

isc_dsql_prepare(
    status_vector,
    &trans, /* Set by previous isc_start_transaction() call. */
    &stmt, /* Statement handle set by this function call. */
    0, /* Specifies statement string is null-terminated. */
    str, /* Statement string. */
    SQLDA_VERSION1, /* XSQLDA version number. */
    out_sqlda /* XSQLDA for storing column data. */
);

```

3. Set up an *XSQLVAR* structure for each column. Setting up an *XSQLVAR* structure involves the following steps:

For columns whose types are known at compile time:

- Specify the column's datatype (if it was not set by *isc_dsql_prepare()*, as previously described).
- Point the *sqldata* field of the *XSQLVAR* to an appropriate local variable.

For columns whose types are not known until run time:

- Coerce the item's datatype (optional), for example, from SQL_VARYING to SQL_TEXT.
- Dynamically allocate local storage for the data pointed to by the *sqldata* field of the *XSQLVAR*.

For both:

- Specify the number of bytes of data to be retrieved into *sqldata*.
- Provide a NULL value indicator for the parameter.

Data retrieval for Blob (and array) columns is different from other types of columns, so the *XSQLVAR* fields must be set differently. For non-Blob (and non-array) columns, *isc_dsql_prepare()* sets each *XSQLVAR sqltype* field to the appropriate field type, and the data retrieved when a select-list row's data is fetched is placed into the *sqldata* space allocated for the column. For Blob columns, the type must be set to SQL_Blob (or SQL_Blob + 1 if a NULL indicator is desired). InterBase stores the internal Blob identifier (Blob ID), not the Blob data, in the *sqldata* space when a row's data is fetched, so you must point *sqldata* to an area the size of a Blob ID. To see how to retrieve the actual Blob data once you have a Blob ID, see **“Reading data from a Blob” on page 121**.

The following code example illustrates the assignments for Blob and non-Blob columns whose types are known at compile time. For examples of handling datatypes that are unknown until run time, see **Chapter 6, “Working with Dynamic SQL.”**

```
#define PROJLEN 20
#define TYPELEN 12
ISC_QUAD blob_id;
char proj_name[PROJLEN + 1];
char prod_type[TYPELEN + 1];
short flag0, flag1, flag2;
out_sqlda->sqlvar[0].sqldata = proj_name;
out_sqlda->sqlvar[0].sqltype = SQL_TEXT + 1;
out_sqlda->sqlvar[0].sqllen = PROJLEN;
out_sqlda->sqlvar[0].sqlind = &flag0;
out_sqlda->sqlvar[1].sqldata = (char *) &blob_id;
out_sqlda->sqlvar[1].sqltype = SQL_Blob + 1;
out_sqlda->sqlvar[1].sqllen = sizeof(ISC_QUAD);
out_sqlda->sqlvar[1].sqlind = &flag1;
out_sqlda->sqlvar[2].sqldata = prod_type;
out_sqlda->sqlvar[2].sqltype = SQL_TEXT + 1;
out_sqlda->sqlvar[2].sqllen = TYPELEN;
out_sqlda->sqlvar[2].sqlind = &flag2;
```

► *Executing the statement*

Once the query statement string is prepared, it can be executed:

```
isc_dsql_execute(
    status_vector,
    &trans,    /* set by previous isc_start_transaction() call */
    &stmt,     /* allocated above by isc_dsql_allocate_statement() */
    1,       /* XSQLDA version number */
    NULL     /* NULL since stmt doesn't have input values */
);
```

This statement creates a select list, the rows returned by execution of the statement.

► *Fetching selected rows*

A looping construct is used to fetch (into the output *XSQLDA*) the column data for a single row at a time from the select-list and to process each row before the next row is fetched. Each execution of *isc_dsql_fetch()* fetches the column data into the corresponding *XSQLVAR* substructures of *out_sqlda*. For the Blob column, the Blob ID, not the actual Blob data, is fetched.

```
ISC_STATUS fetch_stat;
long  SQLCODE;
. . .
while ((fetch_stat =
        isc_dsql_fetch(status_vector, &stmt, 1, out_sqlda))
        == 0)
{
    proj_name[PROJLEN] = '\0';
    prod_type[TYPELEN] = '\0';
    printf("\nPROJECT: %-20s TYPE: %-15s\n\n",
          proj_name, prod_type);
    /* Read and process the Blob data (see next section) */
}
if (fetch_stat != 100L)
{
    /* isc_dsql_fetch returns 100 if no more rows remain to be
       retrieved */
    SQLCODE = isc_sqlcode(status_vector);
    isc_print_sqlerror(SQLCODE, status_vector);
    return(1);
}
```

► *Reading and processing the Blob data*

To read and process the Blob data:

1. Declare and initialize a Blob handle:

```
isc_blob_handle blob_handle; /* Declare a Blob handle. */
blob_handle = NULL; /* Set handle to NULL before using it */
```

2. Create a buffer for holding each Blob segment as it is read. Its size should be the maximum size segment your program expects to be read from the Blob.

```
char blob_segment[80];
```

3. Declare an unsigned short variable into which InterBase will store the actual length of each segment read:

```
unsigned short actual_seg_len;
```

4. Open the Blob with the fetched *blob_id*:

```
isc_open_blob2(
    status_vector,
    &db_handle,
    &trans,
    &blob_handle, /* set by this function to refer to the Blob */
    &blob_id, /* Blob ID put into out_sqlda by isc_dsql_fetch() */
    0, /* BPB length = 0; no filter will be used */
    NULL /* NULL BPB, since no filter will be used */
);
```

5. Read all the Blob data by calling *isc_get_segment()* repeatedly to get each Blob segment and its length. Process each segment read. In the following example, “processing” consists of printing each Blob as it is read:

```
blob_stat = isc_get_segment(
    status_vector,
    &blob_handle, /* set by isc_open_blob2()*/
    &actual_seg_len, /* length of segment read */
    sizeof(blob_segment), /* length of segment buffer */
    blob_segment /* segment buffer */
);
while (blob_stat == 0 || status_vector[1] == isc_segment)
{
    /* isc_get_segment returns 0 if a segment was successfully read.
    */
    /* status_vector[1] is set to isc_segment if only part of a */
    /* segment was read due to the buffer (blob_segment) not being */

```

```

    /* large enough. In that case, the following calls to */
    /* isc_get_segment() read the rest of the buffer. */
    printf("%*.*s", actual_seg_len, actual_seg_len, blob_segment);
    blob_stat = isc_get_segment(status_vector, &blob_handle,
        &actual_seg_len, sizeof(blob_segment), blob_segment);
    printf("\n");
};
printf("\n");

```

6. Close the Blob:

```
isc_close_blob(status_vector, &blob_handle);
```

Writing data to a Blob

Before you can create a new Blob and write data to it, you must do at least one of the following:

- Include Blob data in a row to be inserted into a table.
- Replace the data referenced by a Blob column of a row.
- Update the data referenced by a Blob column of a row.

The entry in a Blob column of a row does not actually contain Blob data. Rather, it has a Blob ID referring to the data, which is stored elsewhere. So, to set or modify a Blob column, you need to set (or reset) the Blob ID stored in it. If a Blob column contains a Blob ID, and you modify the column to refer to a different Blob (or to contain NULL), the Blob referenced by the previously stored Blob ID will be deleted during the next garbage collection.

All these operations require the following steps:

1. Prepare an appropriate DSQL statement. This will be an INSERT statement if you are inserting a new row into a table, or an UPDATE statement for modifying a row. Each of these statements will need a corresponding input *XSQLDA* structure for supplying parameter values to the statement at run time. The Blob ID of a new Blob will be one of the values passed.
2. Create a new Blob, and write data into it.
3. Associate the Blob ID of the new Blob with the Blob column of the table row by executing the UPDATE or INSERT statement.

Note that you cannot update Blob data directly. If you want to modify Blob data, you must:

- Create a new Blob.

- Read the old Blob data into a buffer where you can edit or modify it.
- Write the modified data to the new Blob.
- Prepare and execute an UPDATE statement that will modify the Blob column to contain the Blob ID of the new Blob, replacing the old Blob's Blob ID.

The sections below describe the steps required to insert, replace, or update Blob data.

► *Preparing the UPDATE or INSERT statement*

To prepare an UPDATE or INSERT statement for execution, follow these steps:

1. Elicit an UPDATE or INSERT statement string from the user or create one for inserting a row or updating the row containing the Blob column of interest. For example, the following statement is for updating the Blob column named PROJ_DESC in the row of the table, PROJECT, whose PROJ_ID field contains a value specified at run time:

```
char *upd_str =
    "UPDATE PROJECT SET PROJ_DESC = ? WHERE PROJ_ID = ?";
```

As an example of an INSERT statement, the following inserts a new row containing values in four columns:

```
char *in_str = "INSERT INTO PROJECT (PROJ_NAME, PROJ_DESC, PRODUCT,
    PROJ_ID) VALUES (?, ?, ?, ?)";
```

The remaining steps refer only to UPDATE statements, but the actions apply to INSERT statements as well.

2. Declare a variable to hold the input *XSQLDA* needed to supply parameter values to the UPDATE statement at run time. For example, the following declaration creates an *XSQLDA* called *in_sqlda*:

```
XSQLDA *in_sqlda;
```

3. Allocate memory for the input *XSQLDA* using the *XSQLDA_LENGTH* macro. The *XSQLDA* must contain one *XSQLVAR* substructure for each parameter to be passed to the UPDATE statement. The following statement allocates storage for an input *XSQLDA* (*in_sqlda*) with two *XSQLVAR* substructures:

```
in_sqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(2));
```

4. Set the *version* field of the *XSQLDA* to *SQLDA_VERSION1*, and set the *sqln* field to indicate the number of *XSQLVAR* structures allocated:

```
in_sqlda->version = SQLDA_VERSION1;
in_sqlda->sqln = 2;
```

5. Set up the *XSQLVAR* structure in the *XSQLDA* for each parameter to be passed. Setting up an *XSQLVAR* structure involves the following steps:
 - Specify the item's datatype
 - For parameters whose types are known at compile time: Point the *sqldata* field of the *XSQLVAR* to an appropriate local variable that will contain the data to be passed
 - For parameters whose types are not known until run time: Allocate local storage for the data pointed to by the *sqldata* field of the *XSQLVAR*
 - Specify the number of bytes of data

Data storage for Blob (and array) columns is different from other types of columns, so the *XSQLVAR* fields must be set differently. For non-Blob and non-array columns, input parameter data comes from the space pointed to by *sqldata*. For Blob columns, you must set the type to `SQL_Blob` (or `SQL_Blob + 1` if you want a NULL indicator). Your application must store space for the internal Blob identifier, not the Blob data, in the *sqldata* space. For more information about creating a Blob, storing its ID in the *sqldata* space, and associating the Blob with a column, see **“Creating a new Blob and storing data” on page 128**.

The following code example illustrates the assignments for one text column and one Blob column, where the column types are known at compile time. For examples of handling datatypes that are unknown until run time, see **Chapter 6, “Working with Dynamic SQL.”**

```
#define PROJLEN 5
char proj_id[PROJLEN + 1];
ISC_QUAD blob_id;
in_sqlda->sqlvar[0].sqldata = (char *) &blob_id;
in_sqlda->sqlvar[0].sqltype = SQL_Blob + 1;
in_sqlda->sqlvar[0].sqllen = sizeof(ISC_QUAD);
in_sqlda->sqlvar[1].sqldata = proj_id;
in_sqlda->sqlvar[1].sqltype = SQL_TEXT;
in_sqlda->sqlvar[1].sqllen = 5;
```

The *proj_id* variable should be assigned a value at run time (unless the value is known at compile time). The *blob_id* variable should be set to refer to the newly created Blob, as described in the following sections.

► *Creating a new Blob and storing data*

To create a new Blob containing the data to be written:

1. Declare and initialize a Blob handle:

```
isc_blob_handle blob_handle; /* Declare a Blob handle. */
blob_handle = NULL; /* Set handle to NULL before using it */
```

2. Declare and initialize a Blob ID:

```
ISC_QUAD blob_id; /* Declare a Blob ID. */
blob_id = NULL; /* Set handle to NULL before using it */
```

3. Create a new Blob by calling *isc_create_blob2()*:

```
isc_create_blob2(
    status_vector,
    &db_handle,
    &trans,
    &blob_handle, /* set by this function to refer to the new Blob */
    &blob_id, /* Blob ID set by this function */
    0, /* Blob Parameter Buffer length = 0; no filter will be used
*/
    NULL /* NULL Blob Parameter Buffer, since no filter will be used
*/
);
```

This function creates a new Blob, opens it for write access, and sets *blob_handle* to point to the new Blob.

isc_create_blob2() also assigns the Blob a Blob ID, and sets *blob_id* to point to the Blob ID. Note that *blob_id* is the variable pointed to by the *sqldata* field of the UPDATE statement input parameter that specifies the Blob column to be updated. Thus, when the UPDATE statement is executed, this new Blob will be used to update the Blob column.

4. Write all the data to be written to the Blob by making a series of calls to *isc_put_segment()*. The following example reads lines of data, and concatenates each to the Blob referenced by *blob_handle*. (*get_line()* reads the next line of data to be written.)

```
char *line;
unsigned short len;
. . .
line = get_line();
while (line)
{
    len = strlen(line);
    isc_put_segment(
        status_vector,
```

```

        &blob_handle, /* set by previous isc_create_blob2() */
        len, /* length of buffer containing data to write */
        line /* buffer containing data to write into Blob */
    );
    if (status_vector[0] == 1 && status_vector[1])
    {
        isc_print_status(status_vector);
        return(1);
    };
    line = get_line();
};

```

5. Close the Blob:

```
isc_close_blob(status_vector, &blob_handle);
```

► *Associating the new Blob with the Blob column*

Execute the UPDATE statement to associate the new Blob with the Blob column in the row selected by the statement:

```

isc_dsql_execute_immediate(
    status_vector,
    &db_handle,
    &trans,
    0, /* indicates string to execute is null-terminated */
    upd_str, /* UPDATE statement string to be executed */
    1, /* XSQLDA version number */
    in_sqlda /* XSQLDA supplying parameters to UPDATE statement */
);

```

Deleting a Blob

There are four ways to delete a Blob:

- Delete the row containing the Blob. You can use DSQL to execute a DELETE statement.
- Replace the Blob with a different one. If a Blob column contains a Blob ID, and you modify the column to refer to a different Blob, the Blob referenced by the previously stored Blob ID will be deleted during the next garbage collection.
- Reset to NULL the column referring to the Blob, for example, by using DSQL to execute a statement like the following:

```
UPDATE PROJECT SET PROJ_DESC = NULL WHERE PROJ_ID = 'VBASE'
```

The Blob referenced by the previously stored Blob ID will be deleted during the next garbage collection.

- Discard a Blob after it has been created but before it has been associated with a particular column of a table row. Use the `isc_cancel_blob()` function, as in:

```
isc_cancel_blob(status_vector, &blob_handle);
```

Requesting information about an open Blob

After an application opens a Blob, it can obtain information about the Blob. The `isc_blob_info()` call enables an application to query for Blob information such as the total number of segments in the Blob, or the length, in bytes, of the longest segment.

In addition to a pointer to the error status vector and a Blob handle, `isc_blob_info()` requires two application-provided buffers, an item-list buffer, where the application specifies the information it needs, and a result buffer, where InterBase returns the requested information. An application populates the item-list buffer with information requests prior to calling `isc_blob_info()`, and passes it both a pointer to the item-list buffer, and the size, in bytes, of that buffer.

The application must also create a result buffer large enough to hold the information returned by InterBase. It passes both a pointer to the result buffer, and the size, in bytes, of that buffer to `isc_blob_info()`. If InterBase attempts to pass back more information than can fit in the result buffer, it puts the value, `isc_info_truncated`, defined in `ibase.h`, in the final byte of the result buffer.

Item-list buffer items and result buffer values

The item-list buffer is a char array that holds a sequence of byte values, one per requested item of information. Each byte is an *item type*, specifying the kind of information desired. Compile-time constants for all item types are defined in `ibase.h`:

```
#define isc_info_blob_num_segments      4
#define isc_info_blob_max_segment      5
#define isc_info_blob_total_length     6
#define isc_info_blob_type             7
```

The result buffer returns a series of *clusters* of information, one per item requested. Each cluster consists of three parts:

1. A one-byte *item type*. Each is the same as one of the item types in the item-list buffer.
2. A 2-byte number specifying the number of bytes that follow in the remainder of the cluster.
3. A *value*, stored in a variable number of bytes, whose interpretation depends on the item type.

A calling program is responsible for interpreting the contents of the result buffer and for deciphering each cluster as appropriate.

The clusters returned to the result buffer are not aligned. Furthermore, all numbers are represented in a generic format, with the least significant byte first, and the most significant byte last. Signed numbers have the sign in the last byte. Convert the numbers to a datatype native to your system, if necessary, before interpreting them. The API call, *isc_vax_integer()*, can be used to perform the conversion.

The following table lists items about which information can be requested and returned, and the values reported:

Request and return item	Return value
<i>isc_info_blob_num_segments</i>	Total number of segments
<i>isc_info_blob_max_segment</i>	Length of the longest segment
<i>isc_info_blob_total_length</i>	Total size, in bytes, of Blob
<i>isc_info_blob_type</i>	Type of Blob (0: segmented, or 1: stream)

TABLE 7.2 Blob request and return items

In addition to the information InterBase returns in response to a request, InterBase can also return one or more of the following status messages to the result buffer. Each status message is one unsigned byte in length:

Item	Description
<i>isc_info_end</i>	End of the messages
<i>isc_info_truncated</i>	Result buffer is too small to hold any more requested information
<i>isc_info_error</i>	Requested information is unavailable. Check the status vector for an error code and message

TABLE 7.3 Status message return items

isc_blob_info() call example

The following code requests the number of segments and the maximum segment size for a Blob after it is opened, then examines the result buffer:

```

char blob_items[] = {
    isc_info_blob_max_segment, isc_info_blob_num_segments};
char res_buffer[20], *p, item;
short length;
SLONG max_size = 0L, num_segments = 0L;
ISC_STATUS status_vector[20];
isc_open_blob2(
    status_vector,
    &db_handle, /* database handle, set by isc_attach_database() */
    &tr_handle, /* transaction handle, set by isc_start_transaction()
*/
    &blob_handle, /* set by this function to refer to the Blob */
    &blob_id,     /* Blob ID of the Blob to open */
    0,          /* BPB length = 0; no filter will be used */
    NULL       /* NULL BPB, since no filter will be used */
);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector);
    return(1);
}
isc_blob_info(
    status_vector,
    &blob_handle, /* Set in isc_open_blob2() call above. */
    sizeof(blob_items), /* Length of item-list buffer. */
    blob_items,      /* Item-list buffer. */
    sizeof(res_buffer), /* Length of result buffer. */
    res_buffer       /* Result buffer */
);
if (status_vector[0] == 1 && status_vector[1])
{
    /* An error occurred. */
    isc_print_status(status_vector);
    isc_close_blob(status_vector, &blob_handle);
    return(1);
};
/* Extract the values returned in the result buffer. */

```

```

for (p = res_buffer; *p != isc_info_end ;)
{
    item = *p++
    length = (short)isc_vax_integer(p, 2);
    p += 2;
    switch (item)
    {
        case isc_info_blob_max_segment:
            max_size = isc_vax_integer(p, length);
            break;
        case isc_info_blob_num_segments:
            num_segments = isc_vax_integer(p, length);
            break;
        case isc_info_truncated:
            /* handle error */
            break;
        default:
            break;
    }
    p += length;
};

```

Blob descriptors

A Blob descriptor is used to provide dynamic access to Blob information. For example, it can be used to store information about Blob data for filtering (conversion) purposes, such as character set information for text Blob data and subtype information for text and non-text Blob data. Two Blob descriptors are needed whenever a filter will be used when writing to or reading from a Blob: one to describe the filter source data, and the other to describe the target.

A Blob descriptor is a structure defined in the **ibase.h** header file as follows:

```

typedef struct {
    short blob_desc_subtype; /* type of Blob data */
    short blob_desc_charset; /* character set */
    short blob_desc_segment_size; /* segment size */
    unsigned char blob_desc_field_name [32]; /* Blob column name */
    unsigned char blob_desc_relation_name [32]; /* table name */
} ISC_Blob_DESC;

```

For more information about the character sets recognized by InterBase, see the *Language Reference*.

The segment size of a Blob is the maximum number of bytes that an application is expected to write to or read from the Blob. You can use this size to allocate your own buffers.

The *blob_desc_relation_name* and *blob_desc_field_name* fields contain null-terminated strings.

Populating a Blob descriptor

There are four possible ways to populate a Blob descriptor. You can do so by:

- Calling *isc_blob_default_desc()*. This stores default values into the descriptor fields. The default subtype is 1 (TEXT), segment size is 80 bytes, and charset is the default charset for your process.
- Calling *isc_blob_lookup_desc()*. This accesses the database system metadata tables to look up and copy information for the specified Blob column into the descriptor fields.
- Calling *isc_blob_set_desc()*. This initializes the descriptor from parameters you call it with, rather than accessing the database metadata.
- Setting the descriptor fields directly.

The following example calls *isc_blob_lookup_desc()* to look up the current subtype and character set information for a Blob column named PROJ_DESC in a table named PROJECT. It stores the information into the source descriptor, *from_desc*.

```
isc_blob_lookup_desc (
    status_vector,
    &db_handle; /* Set by previous isc_attach_database() call. */
    &tr_handle, /* Set by previous isc_start_transaction() call. */
    "PROJECT", /* Table name. */
    "PROJ_DESC", /* Column name. */
    &from_desc, /* Blob descriptor filled in by this function call. */
    &global /* Global column name, returned by this function. */
)
```

For more information about the usage of Blob descriptors in applications that request data filtering, and for further examples of populating Blob descriptors, see **“Writing an application that requests filtering” on page 142**.

Filtering Blob data

A Blob filter is a routine that translates Blob data from one subtype to another.

InterBase includes a set of special internal Blob filters that convert from subtype 0 (unstructured data) to subtype 1 (TEXT), and from subtype 1 to subtype 0.

In addition to using these standard filters, you can write your own external filters to provide special data translation. For example, you might develop a filter to convert one image format to another, for instance to display the same image on monitors with different resolutions. Or you might convert a binary Blob to plain text and back again to be able to move the file more easily from one system to another.

If you define filters, you can assign them subtype identifiers from $-32,768$ to -1 .

The following sections provide an overview of how to write Blob filters, followed by details of how to write an application that requires filtering. For more information about writing Blob filters, see the *Embedded SQL Guide*.

Note Blob filters are available for databases residing on all InterBase server platforms except NetWare, where Blob filters cannot be created or used.

Using your own filters

Unlike the standard InterBase filters that convert between subtype 0 and subtype 1, an external Blob filter is generally part of a library of routines you create and link to an application.

You can write Blob filters in C or Pascal (or any language that can be called from C). To use your own filters, follow these steps:

1. Decide which filters you need to write.
2. Write the filters in a host language.
3. Build a shared filter library.
4. Make the filter library available.
5. Define the filters to the database.
6. Write an application that requests filtering.

Steps numbered 2, 5, and 6 are described in greater detail in the following sections.

Declaring an external Blob filter to the database

To declare an external filter to a database, use the `DECLARE FILTER` statement. For example, the following statement declares the filter, `SAMPLE`:

```
DECLARE FILTER SAMPLE
    INPUT_TYPE -1 OUTPUT_TYPE -2
    ENTRY_POINT 'FilterFunction'
    MODULE_NAME 'filter.dll';
```

In the example, the filter's input subtype is defined as `-1` and its output subtype as `-2`. If subtype `-1` specifies lowercase text, and subtype `-2` uppercase text, then the purpose of filter `SAMPLE` would be to translate Blob data from lowercase text to uppercase text.

The `ENTRY_POINT` and `MODULE_NAME` parameters specify the external routine that InterBase calls when the filter is invoked. The `MODULE_NAME` parameter specifies `filter.dll`, the dynamic link library containing the filter's executable code. The `ENTRY_POINT` parameter specifies the entry point into the DLL. Although the example shows only a simple file name, it is good practice to specify a fully-qualified path name, since users of your application need to load the file.

Writing an external Blob filter

If you choose to write your own filters, you must have a detailed understanding of the datatypes you plan to translate. InterBase does not do strict datatype checking on Blob data; it is your responsibility.

► *Defining the filter function*

When writing a filter, you must include an entry point, known as a *filter function*, in the declaration section of the program. InterBase calls the filter function when an application performs a Blob access operation on a Blob specified to use the filter. All communication between InterBase and the filter is through the filter function. The filter function itself may call other functions that comprise the filter executable.

You declare the name of the filter function and the name of the filter executable with the `ENTRY_POINT` and `MODULE_NAME` parameters of the `DECLARE FILTER` statement.

A filter function must have the following declaration *calling sequence*:

```
filter_function_name(short action, isc_blob_ctl control);
```

The parameter, *action*, is one of eight possible action macro definitions, and the parameter, *control*, is an instance of the *isc_blob_ctl* Blob control structure, defined in the InterBase header file, **ibase.h**. These parameters are discussed later in this chapter.

The following listing of a skeleton filter declares the filter function, *jpeg_filter*:

```
#include <ibase.h>
#define SUCCESS 0
#define FAILURE 1
ISC_STATUS jpeg_filter(short action, isc_blob_ctl control)
{
    ISC_STATUS status = SUCCESS;
    switch (action)
    {
        case isc_blob_filter_open:
            . . .
            break;
        case isc_blob_filter_get_segment:
            . . .
            break;
        case isc_blob_filter_create:
            . . .
            break;
        case isc_blob_filter_put_segment:
            . . .
            break;
        case isc_blob_filter_close:
            . . .
            break;
        case isc_blob_filter_alloc:
            . . .
            break;
        case isc_blob_filter_free:
            . . .
            break;
        case isc_blob_filter_seek:
            . . .
            break;
        default:
            . . .
            break;
    }
}
```

```
return status;
}
```

InterBase passes one of eight possible actions to the filter function, *jpeg_filter*, by way of the *action* parameter, and also passes an instance of the Blob control structure, *isc_blob_ctl*, by way of the parameter, *control*.

The ellipses (...) in the previous listing represent code that performs some operations based on each action, or event, that is listed in the case statement. Most of the actions correspond to API functions called by an application. For more information regarding the types of code to write for each action, see the *Embedded SQL Guide*.

► *Defining the Blob control structure*

The *isc_blob_ctl* Blob control structure provides the fundamental method of data exchange between InterBase and a filter.

The Blob control structure is defined as a typedef, *isc_blob_ctl*, in **ibase.h**, as follows:

```
typedef struct isc_blob_ctl {
    ISC_STATUS (*ctl_source)();
    /* Internal InterBase Blob access routine. */
    struct isc_blob_ctl *ctl_source_handle;
    /* Instance of isc_blob_ctl to pass to
       internal InterBase Blob access routine. */
    short ctl_to_sub_type; /* Target subtype. */
    short ctl_from_sub_type; /* Source subtype. */
    unsigned short ctl_buffer_length; /* Length of ctl_buffer. */
    unsigned short ctl_segment_length; /* Length of current segment. */
    unsigned short ctl_bpb_length; /* Blob parameter buffer length. */
    char *ctl_bpb; /* Pointer to Blob parameter buffer. */
    unsigned char *ctl_buffer; /* Pointer to segment buffer. */
    ISC_LONG ctl_max_segment; /* Length of longest Blob segment. */
    ISC_LONG ctl_number_segments; /* Total number of segments. */
    ISC_LONG ctl_total_length; /* Total length of Blob. */
    ISC_STATUS *ctl_status; /* Pointer to status vector. */
    long ctl_data[8]; /* Application-specific data. */
} *ISC_Blob_CTL;
```

The purpose of certain *isc_blob_ctl* fields depend on the action being performed.

For example, when an application calls the *isc_put_segment()* API function, InterBase passes an *isc_blob_filter_put_segment* action to the filter function. The buffer pointed to by the *ctl_buffer* field of the control structure passed to the filter function contains the segment data to be written, as specified by the application in its call to *isc_put_segment()*. Because the buffer contains information passed into the filter function, it is called an IN

field. The filter function should include instructions in the case statement under the *isc_blob_filter_put_segment* case for performing the filtering and then passing the data on for writing to the database. This can be done by calling the **ctl_source* internal InterBase Blob access routine. For more information about *ctl_source*, see the *Embedded SQL Guide*.

On the other hand, when an application calls the *isc_get_segment()* API function, the buffer pointed to by *ctl_buffer* in the control structure passed to a filter function is empty. In this situation, InterBase passes an *isc_blob_filter_get_segment* action to the filter function. The filter function *isc_blob_filter_get_segment* action handling should include instructions for filling *ctl_buffer* with segment data from the database to return to the application. This can be done by calling the **ctl_source* internal InterBase Blob access routine. In this case, because the buffer is used for filter function output, it is called an OUT field.

The following table describes each of the fields in the *isc_blob_ctl* Blob control structure, and whether they are used for filter function input (IN), or output (OUT).

Field name	Description
<i>(*ctl_source)()</i>	Pointer to the internal InterBase Blob access routine (IN)
<i>*ctl_source_handle</i>	Pointer to an instance of <i>isc_blob_ctl</i> to be passed to the internal InterBase Blob access routine (IN)
<i>ctl_to_sub_type</i>	Target subtype: information field provided to support multi-purpose filters that can perform more than one kind of translation; this field and the next one enable such a filter to decide which translation to perform (IN)
<i>ctl_from_sub_type</i>	Source subtype: information field provided to support multi-purpose filters that can perform more than one kind of translation; this field and the previous one enable such a filter to decide which translation to perform (IN)
<i>ctl_buffer_length</i>	For <i>isc_blob_filter_put_segment</i> , field is an IN field that contains the length of the segment data contained in <i>ctl_buffer</i> For <i>isc_blob_filter_get_segment</i> , field is an IN field set to the size of the buffer pointed at by <i>ctl_buffer</i> , which is used to store the retrieved Blob data
<i>ctl_segment_length</i>	Length of current segment. For <i>isc_blob_filter_put_segment</i> , field is not used For <i>isc_blob_filter_get_segment</i> , field is an OUT field set to the size of the retrieved segment (or partial segment, in the case when the buffer length <i>ctl_buffer_length</i> is less than the actual segment length)
<i>ctl_bpb_length</i>	Length of the Blob parameter buffer

TABLE 7.4 *isc_blob_ctl* structure field descriptions

Field name	Description
<i>*ctl_bpb</i>	Pointer to the Blob parameter buffer
<i>*ctl_buffer</i>	Pointer to segment buffer. For <i>isc_blob_filter_put_segment</i> , field is an IN field that contains the segment data For <i>isc_blob_filter_get_segment</i> , field is an OUT field the filter function fills with segment data for return to the application
<i>ctl_max_segment</i>	Length, in bytes, of the longest segment in the Blob. Initial value is 0. The filter function sets this field. This field is information only.
<i>ctl_number_segments</i>	Total number of segments in the Blob. Initial value is 0. The filter function sets this field. This field is information only.
<i>ctl_total_length</i>	Total length, in bytes, of the Blob. Initial value is 0. The filter function sets this field. This field is information only.
<i>*ctl_status</i>	Pointer to InterBase status vector. (OUT)
<i>ctl_data [8]</i>	8-element array of application-specific data. Use this field to store resource pointers, such as memory pointers and file handles created by the <i>isc_blob_filter_open</i> handler, for example. Then, the next time the filter function is called, the resource pointers will be available for use. (IN/OUT)

TABLE 7.4 *isc_blob_ctl* structure field descriptions (continued)

► *Programming filter function actions*

When an application invokes a Blob API function on a Blob to be filtered, InterBase passes a corresponding action message to the filter function by way of the *action* parameter. There are eight possible actions. The following action macro definitions are declared in the **ibase.h** file:

```
#define isc_blob_filter_open 0
#define isc_blob_filter_get_segment 1
#define isc_blob_filter_close 2
#define isc_blob_filter_create 3
#define isc_blob_filter_put_segment 4
#define isc_blob_filter_alloc 5
#define isc_blob_filter_free 6
#define isc_blob_filter_seek 7
```

The following table lists the actions, and specifies when the filter function is invoked with each particular action. Most of the actions are the result of events that occur when an application invokes a Blob API function.

Action	When filter is invoked with corresponding action
<i>isc_blob_filter_open</i>	Invoked when an application calls <i>isc_open_blob2()</i>
<i>isc_blob_filter_get_segment</i>	Invoked when an application calls <i>isc_get_segment()</i>
<i>isc_blob_filter_close</i>	Invoked when an application calls <i>isc_close_blob()</i>
<i>isc_blob_filter_create</i>	Invoked when an application calls <i>isc_create_blob2()</i>
<i>isc_blob_filter_put_segment</i>	Invoked when an application calls <i>isc_put_segment()</i>
<i>isc_blob_filter_alloc</i>	Invoked when InterBase initializes filter processing; not a result of a particular application action
<i>isc_blob_filter_free</i>	Invoked when InterBase ends filter processing; not a result of a particular application action
<i>isc_blob_filter_seek</i>	Reserved for internal filter use; not used by external filters

TABLE 7.5 Action constants

This concludes the overview of writing Blob filters. For detailed information about filters and how to program filter function actions, as well as a reference to a filter application example, see the *Embedded SQL Guide*.

Writing an application that requests filtering

To request filtering of Blob data as it is read from or written to a Blob, follow these steps in your application:

1. Create a Blob parameter buffer (BPB) specifying the source and target subtypes, and optionally character sets (for TEXT subtypes).
2. Call either *isc_open_blob2()* or *isc_create_blob2()* to open a Blob for read or write access, respectively. In the call, pass the BPB, whose information InterBase will use to determine which filter should be called.

► *Understanding the Blob parameter buffer*

A *Blob parameter buffer* (BPB) is needed whenever a filter will be used when writing to or reading from a Blob.

The BPB is a char array variable, specifically declared in an application, that contains the source and target subtypes. When data is read from or written to the Blob associated with the BPB, InterBase will automatically invoke an appropriate filter, based on the source and target subtypes specified in the BPB.

If the source and target subtypes are both 1 (TEXT), and the BPB also specifies different source and target character sets, then when data is read from or written to the Blob associated with the BPB, InterBase will automatically convert each character from the source to the target character set.

A Blob parameter buffer can be generated in one of two ways:

1. Indirectly, through API calls to create source and target descriptors and then generate the BPB from the information in the descriptors.
2. Directly by populating the BPB array with appropriate values.

If you generate a BPB via API calls, you do not need to know the format of the BPB. But if you wish to directly generate a BPB, then you must know the format.

Both approaches are described in the following sections. The format of the BPB is documented in the section about directly populating the BPB.

GENERATING A BLOB PARAMETER BUFFER USING API CALLS

To generate a BPB indirectly, use API calls to create source and target Blob descriptors, and then call *isc_blob_gen_bpb()* to generate the BPB from the information in the descriptors. Follow these steps:

1. Declare two Blob descriptors, one for the source, and the other for the target. For example,

```
#include <ibase.h>
ISC_Blob_DESC from_desc, to_desc;
```

2. Store appropriate information in the Blob descriptors, by calling one of the functions *isc_blob_default_desc()*, *isc_blob_lookup_desc()*, or *isc_blob_set_desc()*, or by setting the descriptor fields directly. The following example looks up the current subtype and character set information for a Blob column named GUIDEBOOK in a table named TOURISM, and stores it into the source descriptor, *from_desc*. It then sets the target descriptor, *to_desc* to the default subtype (TEXT) and character set, so that the source data will be converted to plain text.

```
isc_blob_lookup_desc (
    status_vector,
    &db_handle; /* set in previous isc_attach_database() call */
    &tr_handle, /* set in previous isc_start_transaction() call */
    "TOURISM", /* table name */
```

```

    "GUIDEBOOK", /* column name */
    &from_desc, /* Blob descriptor filled in by this function call */
    &global);
if (status_vector[0] == 1 && status_vector[1])
{
    /* process error */
    isc_print_status(status_vector);
    return(1);
};
isc_blob_default_desc (
    &to_desc, /* Blob descriptor filled in by this function call */
    "", /* NULL table name; it's not needed in this case */
    ""); /* NULL column name; it's not needed in this case */

```

For more information about Blob descriptors, see **“Blob descriptors” on page 134**.

3. Declare a character array which will be used as the BPB. Make sure it is at least as large as all the information that will be stored in the buffer.

```
char bpb[20];
```

4. Declare an unsigned short variable into which InterBase will store the actual length of the BPB data:

```
unsigned short actual_bpb_length;
```

5. Call *isc_blob_gen_bpb()* to populate the BPB based on the source and target Blob descriptors passed to *isc_blob_gen_bpb()*. For example,

```

isc_blob_gen_bpb(
    status_vector,
    &to_desc, /* target Blob descriptor */
    &from_desc, /* source Blob descriptor */
    sizeof(bpb), /* length of BPB buffer */
    bpb, /* buffer into which the generated BPB will be stored */
    &actual_bpb_length /* actual length of generated BPB */
);

```

GENERATING A BLOB PARAMETER BUFFER DIRECTLY

It is possible to generate a BPB directly.

A BPB consists of the following parts:

1. A byte specifying the version of the parameter buffer, always the compile-time constant, *isc_bpb_version1*.
2. A contiguous series of one or more *clusters* of bytes, each describing a single parameter.

Each cluster consists of the following parts:

1. A one-byte parameter type. There are compile-time constants defined for all the parameter types (for example, *isc_bpb_target_type*).
2. A one-byte number specifying the number of bytes that follow in the remainder of the cluster.
3. A variable number of bytes, whose interpretation depends on the parameter type.

Note All numbers in the Blob parameter buffer must be represented in a generic format, with the least significant byte first, and the most significant byte last. Signed numbers should have the sign in the last byte. The API function *isc_vax_integer()* can be used to reverse the byte order of a number. For more information about *isc_vax_integer()*, see **“isc_vax_integer()” on page 391**.

The following table lists the parameter types and their meaning:

Parameter type	Description
<i>isc_bpb_target_type</i>	Target subtype
<i>isc_bpb_source_type</i>	Source subtype
<i>isc_bpb_target_interp</i>	Target character set
<i>isc_bpb_source_interp</i>	Source character set

TABLE 7.6 Blob parameter buffer parameter types

The BPB must contain *isc_bpb_version1* at the beginning, and must contain clusters specifying the source and target subtypes. Character set clusters are optional. If the source and target subtypes are both 1 (TEXT), and the BPB also specifies different source and target character sets, then when data is read from or written to the Blob associated with the BPB, InterBase will automatically convert each character from the source to the target character set.

The following is an example of directly creating a BPB for a filter whose source subtype is `-4` and target subtype is `1` (TEXT):

```
char bpb[] = {
    isc_bpb_version1,
    isc_bpb_target_type,
    1,      /* # bytes that follow which specify target subtype */
    1,      /* target subtype (TEXT) */
    isc_bpb_source_type,
    1,      /* # bytes that follow which specify source subtype */
    -4,    /* source subtype*/
};
```

Of course, if you do not know the source and target subtypes until run time, you can assign those values in the appropriate BPB locations at run time.

► *Requesting filter usage*

You request usage of a filter when opening or creating a Blob for read or write access. In the call to `isc_open_blob2()` or `isc_create_blob2()`, pass the BPB, whose information InterBase will use to determine which filter should be called.

The following example illustrates creating and opening a Blob for write access. For further information about writing data to a Blob and updating a Blob column of a table row to refer to the new Blob, see **“Writing data to a Blob” on page 126**.

Opening a Blob for read access requires additional steps to select the appropriate Blob to be opened. For more information, see **“Reading data from a Blob” on page 121**.

```
isc_blob_handle blob_handle; /* declare at beginning */
ISC_QUAD blob_id; /* declare at beginning */
. . .
isc_create_blob2(
    status_vector,
    &db_handle,
    &tr_handle,
    &blob_handle, /* to be filled in by this function */
    &blob_id,     /* to be filled in by this function */
    actual_bpb_length, /* length of BPB data */
    &bpb         /* Blob parameter buffer */
)
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector);
    return(1);
}
```

}

Working with Array Data

This chapter describes arrays of datatypes and how to work with them using API functions. It shows how to set up an array descriptor specifying the array or array subset to be retrieved or written to, and how to use the two API functions that control access to arrays.

The following table summarizes the API functions for working with arrays. First the functions that can be used to populate an array descriptor are listed, followed by the functions for accessing array data.

Function	Purpose
<i>isc_array_lookup_desc()</i>	Looks up and stores into an array descriptor the datatype, length, scale, and dimensions for all elements in the specified array column of the specified table
<i>isc_array_lookup_bounds()</i>	Performs the same actions as the function, <i>isc_array_lookup_desc()</i> , but also looks up and stores the upper and lower bounds of each dimension
<i>isc_array_set_desc()</i>	Initializes an array descriptor from parameters passed to it
<i>isc_array_get_slice()</i>	Retrieves data from an array
<i>isc_array_put_slice()</i>	Writes data to an array

TABLE 8.1 API array access functions

Introduction to arrays

InterBase supports arrays of most datatypes. Using an array enables multiple data items to be stored in a single column. InterBase can treat an array as a single unit, or as a series of separate units, called *slices*. Using an array is appropriate when:

- The data items naturally form a set of the same datatype.
- The entire set of data items in a single database column must be represented and controlled as a unit, as opposed to storing each item in a separate column.
- Each item must also be identified and accessed individually.

The data items in an array are called *array elements*. An array can contain elements of any InterBase datatype except Blob, and cannot be an array of arrays. All of the elements of a particular array are of the same datatype.

InterBase supports multi-dimensional arrays, arrays with 1 to 16 dimensions. Multi-dimensional arrays are stored in row-major order.

Array dimensions have a specific range of upper and lower boundaries, called *subscripts*. The array subscripts are defined when an array column is created. For information about creating an array, see the *Language Reference*.

Array database storage

InterBase does not store array data directly in the array field of a table record. Instead, it stores an *array ID* there. The array ID is a unique numeric value that references the array data, which is stored elsewhere in the database.

Array descriptors

An *array descriptor* describes an array or array subset to be retrieved or written to the `ISC_ARRAY_DESC` structure. `ISC_ARRAY_DESC` is defined in the InterBase `ibase.h` header file as follows:

```
typedef struct {
    unsigned char array_desc_dtype; /* Datatype */
    char array_desc_scale; /* Scale for numeric datatypes */
    unsigned short array_desc_length;
    /* Length in bytes of each array element */
    char array_desc_field_name [32]; /* Column name */
    char array_desc_relation_name [32]; /* Table name */
    short array_desc_dimensions; /* Number of array dimensions */
    short array_desc_flags;
    /* Specifies whether array is to be accessed in row-major or
       column-major order */
    ISC_ARRAY_BOUND array_desc_bounds [16];
    /* Lower and upper bounds for each dimension */
} ISC_ARRAY_DESC;
```

`ISC_ARRAY_BOUND` is defined as:

```
typedef struct {
    short array_bound_lower; /* lower bound */
    short array_bound_upper; /* upper bound */
} ISC_ARRAY_BOUND;
```

An array descriptor contains 16 `ISC_ARRAY_BOUND` structures, one for each possible dimension. An array with n dimensions has upper and lower bounds set for the first n `ISC_ARRAY_BOUND` structures. The number of actual array dimensions is specified in the `array_desc_dimensions` field of the array descriptor.

When you retrieve data from an array, you supply an array descriptor defining the array *slice* (entire array or subset of contiguous array elements) to be retrieved. Similarly, when you write data to an array, you supply an array descriptor defining the array slice to be written to.

Populating an array descriptor

There are four ways to populate an array descriptor:

- Call `isc_array_lookup_desc()`, which looks up (in the system metadata tables) and stores in an array descriptor the datatype, length, scale, and dimensions for a specified array column in a specified table. This function also stores the table and column name in the descriptor, and initializes its `array_desc_flags` field to indicate that the array is to be accessed in row-major order. For example,

```
isc_array_lookup_desc(
    status_vector,
    &db_handle,      /* Set by isc_attach_database() */
    &tr_handle,     /* Set by isc_start_transaction() */
    "PROJ_DEPT_BUDGET", /* table name */
    "QUART_HEAD_CNT", /* array column name */
    &desc          /* descriptor to be filled in */
);
```

- Call `isc_array_lookup_bounds()`, which looks and functions the same as a call to `isc_array_lookup_desc()`, except that the function `isc_array_lookup_bounds()` also looks up and stores into the array descriptor the upper and lower bounds of each dimension.
- Call `isc_array_set_desc()`, which initializes the descriptor from parameters, rather than by accessing the database metadata. For example,

```
short dtype = SQL_TEXT;
short len = 8;
short numdims = 2;
isc_array_set_desc(
    status_vector,
    "TABLE1",      /* table name */
    "CHAR_ARRAY", /* array column name */
    &dtype,        /* datatype of elements */
    &len,          /* length of each element */
    &numdims,     /* number of array dimensions */
    &desc         /* descriptor to be filled in */
);
```

- Setting the descriptor fields directly. An example of setting the `array_desc_dimensions` field of the descriptor, `desc`, is:

```
desc.array_desc_dimensions = 2;
```

For complete syntax and information about *isc_array_lookup_bounds()*, *isc_array_lookup_desc()*, and *isc_array_set_desc()*, see **Chapter 13, “API Function Reference.”**

Accessing array data

InterBase supports the following operations on array data:

- Reading from an array or array slice.
- Writing to an array:
 - Including a new array in a row to be inserted into a table.
 - Replacing the array referenced by an array column of a row with a new array.
 - Updating the array referenced by an array column of a row by modifying the array data or a slice of the data.
- Deleting an array.

Dynamic SQL (DSQL) API functions and the *XSQLDA* data structure are needed to execute SELECT, INSERT, and UPDATE statements required to select, insert, or update relevant array data. The following sections include descriptions of the DSQL programming methods required to execute the sample statements provided.

For more information about DSQL and the *XSQLDA*, see **Chapter 6, “Working with Dynamic SQL.”**

Note The following array operations are *not* supported:

- Referencing array dimensions dynamically in DSQL.
- Setting individual array elements to NULL.
- Using aggregate functions, such as MIN() and MAX(), with arrays.
- Referencing arrays in the GROUP BY clause of a SELECT.
- Creating views that select from array slices.

Reading data from an array

There are seven steps required for reading data from an array or slice of an array:

1. Create a `SELECT` statement that specifies selection of the array column (and any other columns desired) in the rows of interest.
2. Prepare an output `XSQLDA` structure to hold the column data for each row that is fetched.
3. Prepare the `SELECT` statement for execution.
4. Execute the statement.
5. Populate an array descriptor with information describing the array or array slice to be retrieved.
6. Fetch the selected rows one by one.
7. Read and process the array data from each row.

► *Creating the SELECT statement*

Elicit a statement string from the user or create one that consists of the SQL query that will select rows containing the array data of interest. In your query, specify the array column name and the names of any other columns containing data you are interested in. For example, the following creates an SQL query statement string that selects an array column named `QUART_HEAD_CNT` and another column named `DEPT_NO` from the table, `PROJ_DEPT_BUDGET`:

```
char *sel_str =
    "SELECT DEPT_NO, QUART_HEAD_CNT FROM PROJ_DEPT_BUDGET \
    WHERE year = 1994 AND PROJ_ID = 'VBASE'";
```

► *Preparing the output XSQLDA*

Most queries return one or more rows of data, referred to as a select-list. An output `XSQLDA` must be created to store the column data for each row that is fetched. For an array column, the column data is an internal array identifier (array ID) that is needed to access the actual data. To prepare the `XSQLDA`, follow these steps:

1. Declare a variable to hold the `XSQLDA`. For example, the following declaration creates an `XSQLDA` called `out_sqlda`:

```
XSQLDA *out_sqlda;
```

2. Allocate memory for the *XSQLDA* using the *XSQLDA_LENGTH* macro. The *XSQLDA* must contain one *XSQLVAR* substructure for each column to be fetched. The following statement allocates storage for an output *XSQLDA* (*out_sqllda*) with two *XSQLVAR* substructures:

```
out_sqllda = (XSQLDA *)malloc(XSQLDA_LENGTH(2));
```

3. Set the *version* field of the *XSQLDA* to *SQLDA_VERSION1*, and set the *sqln* field of the *XSQLDA* to indicate the number of *XSQLVAR* substructures allocated:

```
out_sqllda->version = SQLDA_VERSION1;
out_sqllda->sqln = 2;
```

► *Preparing the SELECT statement for execution*

After an *XSQLDA* is created for holding the column data for each selected row, the query statement string can be prepared for execution. Follow these steps:

1. Declare and initialize an SQL statement handle, then allocate it with *isc_dsql_allocate_statement()*:

```
isc_stmt_handle stmt; /* Declare a statement handle. */
stmt = NULL;          /* Set handle to NULL before allocation. */
isc_dsql_allocate_statement(status_vector, &db_handle, &stmt);
```

2. Ready the statement string for execution with *isc_dsql_prepare()*. This checks the string (*sel_str*) for syntax errors, parses it into a format that can be efficiently executed, and sets the statement handle (*stmt*) to refer to this parsed format. The statement handle is used in a later call to *isc_dsql_execute()*.

If *isc_dsql_prepare()* is passed a pointer to the output *XSQLDA*, as in the following example, it will fill in most fields of the *XSQLDA* and all its *XSQLVAR* substructures with information such as the datatype, length, and name of the corresponding columns in the statement.

A sample call to *isc_dsql_prepare()* is:

```
isc_dsql_prepare(
    status_vector,
    &trans,      /* Set by previous isc_start_transaction() call. */
    &stmt,      /* Statement handle set by this function call. */
    0,         /* Specifies statement string is null-terminated. */
    sel_str,   /* Statement string. */
    1,        /* XSQLDA version number. */
    out_sqllda /* XSQLDA for storing column data. */
);
```

3. Set up an *XSQLVAR* structure for each column. Setting up an *XSQLVAR* structure involves the following steps:

For columns whose types are known at compile time:

- Specify the column's datatype (if it was not set by *isc_dsql_prepare()*, as previously described).
- Point the *sqldata* field of the *XSQLVAR* to an appropriate local variable.

For columns whose types are not known until run time:

- Coerce the item's datatype (optional); for example, from *SQL_VARYING* to *SQL_TEXT*.
- Dynamically allocate local storage for the data pointed to by the *sqldata* field of the *XSQLVAR*.

For both:

Provide a NULL value indicator for the parameter.

- Data retrieval for array (and Blob) columns is different from other types of columns, so the *XSQLVAR* fields must be set differently. For non-array (and non-Blob) columns, *isc_dsql_prepare()* sets each *XSQLVAR* *sqltype* field to the appropriate field type, and the data retrieved when a select list row's data is fetched is placed into the *sqldata* space allocated for the column. For array columns, the type is set to *SQL_ARRAY* (or *SQL_ARRAY + 1* if the array column is allowed to be NULL). InterBase stores the internal array identifier (array ID), not the array data, in the *sqldata* space when a row's data is fetched, so you must point *sqldata* to an area the size of an array ID. To see how to retrieve the actual array or array slice data once you have an array ID, see **“Reading and processing the array data” on page 158**.
- The following code example illustrates the assignments for array and non-array columns whose types are known at compile time. For more information about DSQL and the *XSQLDA*, and working with columns whose types are unknown until run time, see **Chapter 6, “Working with Dynamic SQL.”**

```
ISC_QUAD array_id = 0L;
char dept_no[6];
short flag0, flag1;
out_sqlda->sqlvar[0].sqldata = (char *) dept_no;
out_sqlda->sqlvar[0].sqltype = SQL_TEXT + 1;
out_sqlda->sqlvar[0].sqlind = &flag0;
out_sqlda->sqlvar[1].sqldata = (char *) &array_id;
out_sqlda->sqlvar[1].sqltype = SQL_ARRAY + 1;
out_sqlda->sqlvar[1].sqlind = &flag1;
```

► *Executing the statement*

Once the query statement string is prepared, it can be executed:

```
isc_dsqli_execute(
    status_vector,
    &trans, /* set by previous isc_start_transaction() call */
    &stmt, /* set above by isc_dsqli_prepare() */
    1, /* XSQLDA version number */
    NULL /* NULL since stmt doesn't have input values */
);
```

This statement creates a select-list, the rows returned by execution of the statement.

► *Populating the array descriptor*

To prepare an array descriptor that describes the array or array slice to be read, follow these steps:

1. Create the array descriptor:

```
ISC_ARRAY_DESC desc;
```

2. Fill in the descriptor with information regarding the array column from which data will be read. Do this either by calling one of the functions *isc_array_lookup_bounds()*, *isc_array_lookup_desc()*, or *isc_array_set_desc()*, or by directly filling in the descriptor. For information on the contents of array descriptors, and the usage of these functions, see **“Array descriptors” on page 151**.

Ensure the descriptor boundaries are set to those of the slice to be read.

If you want to retrieve all the array data (that is, not just a smaller slice), set the boundaries to the full boundaries as initially declared for the array column. This is guaranteed to be the case if you fill in the descriptor by calling *isc_array_lookup_bounds()*, as in:

```
ISC_ARRAY_DESC desc;
isc_array_lookup_bounds(
    status_vector,
    &db_handle,
    &trans,
    "PROJ_DEPT_BUDGET", /* table name */
    "QUART_HEAD_CNT", /* array column name */
    &desc);
```

Suppose the array column, `QUART_HEAD_CNT`, is a one-dimensional array consisting of four elements, and it was declared to have a lower subscript bound of 1 and an upper bound of 4 when it was created. Then after the above call to `isc_array_lookup_bounds()`, the array descriptor fields for the boundaries contain the following information:

```
desc.array_desc_bounds[0].array_bound_lower == 1
desc.array_desc_bounds[0].array_bound_upper == 4
```

If you want to read just a slice of the array, then modify the upper and/or lower bounds appropriately. For example, if you just want to read the first two elements of the array, then modify the upper bound to the value 2, as in:

```
desc.array_desc_bounds[0].array_bound_upper = 2
```

► *Fetching selected rows*

A looping construct is used to fetch (into the output `XSQLDA`) the column data for a single row at a time from the select-list and to process each row before the next row is fetched. Each execution of `isc_dsqli_fetch()` fetches the column data for the next row into the corresponding `XSQLVAR` structures of `out_sqlda`. For the array column, the array ID, not the actual array data, is fetched.

```
ISC_STATUS fetch_stat;
long SQLCODE;
. . .
while ((fetch_stat = j
        isc_dsqli_fetch(status_vector, &stmt, 1, out_sqlda))
        == 0)
{
    /* Read and process the array data */
}
if (fetch_stat != 100L)
{
    /* isc_dsqli_fetch returns 100 if no more rows remain to be
       retrieved */
    SQLCODE = isc_sqlcode(status_vector);
    isc_print_sqlerror(SQLCODE, status_vector);
    return(1);
}
```

► *Reading and processing the array data*

To read and process the array or array slice data:

1. Create a buffer for holding the array data to be read. Make it large enough to hold all the elements in the slice to be read (which could be the entire array). For example, the following declares an array buffer large enough to hold 4 long elements:

```
long hcnt[4];
```

2. Declare a short variable for specifying the size of the array buffer:

```
short len;
```

3. Set the variable to the buffer length:

```
len = sizeof(hcnt);
```

4. Read the array or array slice data into the buffer by calling *isc_array_get_slice()*. Process the data read. In the following example, the array is read into the *hcnt* array buffer, and “processing” consists of printing the data:

```
isc_array_get_slice(
    status_vector,
    &db_handle, /* set by isc_attach_database() */
    &trans,     /* set by isc_start_transaction() */
    &array_id, /* array ID put into out_sqlda by isc_dsql_fetch() */
    &desc,     /* array descriptor specifying slice to be read */
    (void *) hcnt, /* buffer into which data will be read */
    (long *) &len /* length of buffer */
);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector);
    return(1);
}
/* Make dept_no a null-terminated string */
dept_no[out_sqlda->sqlvar[0].sqlllen] = '\0';
printf("Department #: %s\n\n", dept_no);
printf("\tCurrent head counts: %ld %ld %ld %ld\n",
    hcnt[0], hcnt[1], hcnt[2], hcnt[3]);
```

Writing data to an array

`isc_array_put_slice()` is called to write data to an array or array slice. Use it to:

- Include a new array in a row to be inserted into a table.
- Replace the current contents of an array column of a row with a new array.
- Update the array referenced by an array column of a row by modifying the array data or a slice of the data.

The entry in an array column of a row does not actually contain array data. Rather, it has an array ID referring to the data, which is stored elsewhere. So, to set or modify an array column, you need to set or change the array ID stored in it. If an array column contains an array ID, and you modify the column to refer to a different array (or to contain NULL), the array referenced by the previously stored array ID will be deleted during the next garbage collection.

The following steps are required to insert, replace, or update array data:

1. Prepare an array descriptor with information describing the array (or slice) to be written to.
2. Prepare an array buffer with the data to be written.
3. Prepare an appropriate DSQL statement. This will be an INSERT statement if you are inserting a new row into a table, or an UPDATE statement for modifying an existing row.
4. Call `isc_array_put_slice()` to create a new array (possibly copying an existing one), and to write the data from the array buffer into the array or array slice.
5. Associate the new array with an array column of the table row being modified or inserted by executing the UPDATE or INSERT statement. This sets the array column to contain the array ID of the new array.

► *Preparing the array descriptor*

To prepare an array descriptor that specifies the array or array slice to be written to, follow these steps:

1. Create the array descriptor:

```
ISC_ARRAY_DESC desc;
```

2. Fill in the descriptor with information regarding the array column to which data will be written. Do this either by calling one of the functions `isc_array_lookup_bounds()`, `isc_array_lookup_desc()`, or `isc_array_set_desc()`, or by directly filling in the descriptor. For information on the contents of array descriptors, and the usage of these functions, see **“Array descriptors” on page 151**.

Ensure the descriptor boundaries are set to those of the slice to be written to.

If you want to write to the entire array rather than to just a slice, set the boundaries to the full boundaries as initially declared for the array column. This is guaranteed to be the case if you fill in the descriptor by calling `isc_array_lookup_bounds()`, as in:

```
isc_array_lookup_bounds(
    status_vector,
    db_handle,
    &trans,
    "PROJ_DEPT_BUDGET", /* table name */
    "QUART_HEAD_CNT", /* array column name */
    &desc);
```

Suppose the array column, `QUART_HEAD_CNT`, is a one-dimensional array consisting of four elements, and it was declared to have a lower subscript bound of 1 and an upper bound of 4 when it was created. Then after a call to `isc_array_lookup_bounds()`, the array descriptor fields for the boundaries contain the following information:

```
desc.array_desc_bounds[0].array_bound_lower == 1
desc.array_desc_bounds[0].array_bound_upper == 4
```

If you just want to write to (or modify) a slice of the array, then change the upper and lower bound appropriately. For example, if you just want to write to the first two elements of the array, then modify the upper bound to the value 2, as in:

```
desc.array_desc_bounds[0].array_bound_upper == 2
```

► *Preparing the array buffer with data*

Create an array buffer to hold the data to be written to the array. Make it large enough to hold all the elements in the slice to be written (which could be the entire array). For example, the following declares an array buffer large enough to hold 4 long elements:

```
long hcnt[4];
```

1. Create a variable specifying the length of the array buffer:

```
short len;
len = sizeof(hcnt);
```

2. Fill the array buffer with the data to be written.

If you are creating a new array, then fill the buffer with data. For example,

```
hcnt[0] = 4;
hcnt[1] = 5;
hcnt[2] = 6;
hcnt[3] = 6;
```

To modify existing array data instead of creating a new one, then perform all the steps listed in **“Reading data from an array” on page 154** to read the existing array data into the array buffer. Modify the data in the buffer.

► *Preparing the UPDATE or INSERT statement*

To prepare an UPDATE or INSERT statement for execution, follow these steps:

1. Elicit an UPDATE or INSERT statement string from the user or create one for inserting a new row or updating the row(s) containing the array column(s) of interest. For example, the following statement is for updating the array column named QUART_HEAD_CNT in the specified row of the table, PROJ_DEPT_BUDGET. The department number and quarterly headcounts are assumed to be supplied at run time:

```
char *upd_str =
    "UPDATE PROJ_DEPT_BUDGET SET QUART_HEAD_CNT = ? WHERE \
    YEAR = 1994 AND PROJ_ID = "MKTPR" AND DEPT_NO = ?";
```

As an example of an INSERT statement, the following is for inserting a new row into the PROJ_DEPT_BUDGET table, with column data supplied at run time:

```
char *upd_str =
    "INSERT INTO PROJ_DEPT_BUDGET (YEAR, PROJ_ID, DEPT_NO, \
    QUART_HEAD_CNT) VALUES (?, ?, ?, ?)";
```

The remaining steps refer only to UPDATE statements, but the actions apply to INSERT statements as well.

2. Declare a variable to hold the input *XSQLDA* needed to supply parameter values to the UPDATE statement at run time. For example, the following declaration creates an *XSQLDA* called *in_sqlda*:

```
XSQLDA *in_sqlda;
```

3. Allocate memory for the input *XSQLDA* using the *XSQLDA_LENGTH* macro. The *XSQLDA* must contain one *XSQLVAR* substructure for each parameter to be passed to the UPDATE statement. The following statement allocates storage for an input *XSQLDA* (*in_sqlda*) with two *XSQLVAR* substructures:

```
in_sqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(2));
```

4. Set the Version field of the *XSQLDA* to *SQLDA_VERSION1*, and set the *Sqln* field to indicate the number of *XSQLVAR* structures allocated:

```
in_sqlda->version = SQLDA_VERSION1;
in_sqlda->sqln = 2;
```

5. Set up the *XSQLVAR* structure in the *XSQLDA* for each parameter to be passed. Setting up an *XSQLVAR* structure involves the following steps:

- Specify the item's datatype.
- For parameters whose types are known at compile time, point the *Ssqldata* field of the *XSQLVAR* to an appropriate local variable that will contain the data to be passed.
- For parameters whose types are not known until run time, allocate local storage for the data pointed to by the *Sqldata* field of the *XSQLVAR*.
- Specify the number of bytes of data.

Data storage for array (and Blob) columns is different from other types of columns, so the *XSQLVAR* fields must be set differently. For non-array (and non-Blob) columns, input parameter data comes from the space pointed to by *Sqldata*. For array columns, set the type to *SQL_ARRAY* (or *SQL_ARRAY + 1* if the array column is allowed to be NULL). The application must store space for the internal array identifier, not the array data, in the *Sqldata* space. See the following sections to create or modify an array, store its array ID in the *Sqldata* space, and then associate the actual array data with the column.

The following code example illustrates the assignments for one TEXT column and one array column, where the column types are known at compile time.

```
#define NUMLEN 4
char dept_no[NUMLEN + 1];
ISC_QUAD array_id;
in_sqlda->sqlvar[0].sqldata = &array_id;
in_sqlda->sqlvar[0].sqltype = SQL_ARRAY + 1;
in_sqlda->sqlvar[0].sqln = sizeof(ISC_QUAD);
in_sqlda->sqlvar[1].sqldata = dept_no;
in_sqlda->sqlvar[1].sqltype = SQL_TEXT;
in_sqlda->sqlvar[1].sqln = 4;
```

The *dept_no* variable should be assigned a value at run time (unless the value is known at compile time). The *array_id* variable should be set to refer to the newly created array, as described in the following sections.

For examples of handling data whose types are not known until run time, see **Chapter 6, “Working with Dynamic SQL.”**

► *Calling isc_array_put_slice()*

The following steps are required to store the data into an array or array slice:

1. Declare an array ID:

```
ISC_QUAD array_id; /* Declare an array ID. */
```

2. Initialize the array ID. If you are creating a new array to be inserted into a new row, or to replace an existing array, then simply initialize the array ID to NULL:

```
array_id = NULL; /* Set handle to NULL before using it */
```

If you are modifying an existing array, then follow the steps listed under “Reading Data from an Array” to read the existing array ID into *array_id*.

3. Call *isc_array_put_slice()*. In your call you pass the array ID (either the array ID of an existing array, or NULL for a new array) in the *array_id* variable. You also pass the buffer of data to be written and a descriptor specifying the array slice to which the data belongs.

When *isc_array_put_slice()* is called with an array ID of an existing array, it creates a new array with the same characteristics as the specified array, and copies the existing array data to the new array. Then *isc_array_put_slice()* writes the data from the array buffer to the new array (or slice of the array), per the bounds specified in the array descriptor, and returns in the same *array_id* variable the array ID of the new array.

When *isc_array_put_slice()* is called with a NULL array ID, it creates a new empty array with characteristics as declared for the array column whose name and table name are specified in the array descriptor passed to *isc_array_put_slice()*. It then writes the data from the array buffer to the new array (or slice of the array), and returns in the *array_id* variable the array ID of the new array.

Note that in both cases, a new array is created, and its array ID is returned in the *array_id* variable. The array is temporary until an UPDATE or INSERT statement is executed to associate the array with a particular column of a particular row.

You can make a single call to `isc_array_put_slice()` to write all the data to the array. Or, you may call `isc_array_put_slice()` multiple times to store data into various slices of the array. In this case, each call to `isc_array_put_slice()` after the first call should pass the array ID of the temporary array. When `isc_array_put_slice()` is called with the array ID of a temporary array, it copies the specified data to the specified slice of the temporary array, but does not create a new array.

The following is a sample call to `isc_array_put_slice()`:

```
isc_array_put_slice(
    status_vector,
    &db_handle,
    &trans,
    &array_id, /* array ID (NULL, or existing array's array ID) */
    &desc,     /* array descriptor describing where to write data */
    hcnt,     /* array buffer containing data to write to array */
    &len      /* length of array buffer */
);
```

This call creates a new array, copies the data in `hcnt` to the new array (or slice of the array), assigns the array an array ID, and sets `array_id` to point to the array ID.

IMPORTANT `array_id` should be the variable pointed to by the `Sqldata` field of the UPDATE (or INSERT) statement input parameter that specifies the array column to be updated. Thus, when the INSERT or UPDATE statement is executed, this new array's array ID will be used to set or update the array column to refer to the new array.

► *Associating the new array with the array column*

Execute the UPDATE statement to associate the new array with the array column in the row selected by the statement:

```
isc_dsqli_execute_immediate(
    status_vector,
    &db_handle,
    &trans,
    0,          /* indicates string to execute is null-terminated */
    upd_str,   /* UPDATE statement string to be executed */
    1,        /* XSQLDA version number */
    in_sqlda  /* XSQLDA supplying parameters to UPDATE statement */
);
```

This sets the array column in the row specified in the UPDATE statement to contain the array ID of the new array. The array ID comes from the `array_id` variable pointed to by the `in_sqlda` parameter corresponding to the array column.

If the array column in the specified row contains the array ID of a different array before the UPDATE statement is executed, then the column is modified to contain the new array ID, and the array referenced by the previously stored array ID will be deleted during the next garbage collection.

Deleting an array

There are three ways to delete an array:

1. Delete the row containing the array. You can use DSQL to execute a DELETE statement.
2. Replace the array with a different one, as described above. If an array column contains an array ID, and you modify the column to refer to a different array, the array referenced by the previously stored array ID will be deleted during the next garbage collection.
3. Reset to NULL the column referring to the array. For example, use DSQL to execute a statement like the following, where LANGUAGE_REQ is an array column:

```
"UPDATE JOB SET LANGUAGE_REQ = NULL \  
WHERE JOB_CODE = "SA12" AND JOB_GRADE = 10"
```

The array referenced by the previously stored array ID will be deleted during the next garbage collection.

Working with Conversions

InterBase uses a proprietary format for internal storage of `TIMESTAMP`, `TIME`, and `DATE` data, but provides the following API calls for translating to and from this format:

- `isc_decode_sql_date()` converts the InterBase internal date format to the C *date* structure
- `isc_encode_sql_date()` converts the C *date* structure to the internal InterBase date format
- `isc_decode_sql_time()` converts the InterBase internal time format to the C *time* structure
- `isc_encode_sql_time()` converts the C *time* structure to the internal InterBase time format
- `isc_decode_timestamp()` converts the InterBase internal timestamp format to the C *timestamp* structure; this call was formerly `isc_decode_date()`
- `isc_encode_timestamp()` converts the C *timestamp* structure to the internal InterBase timestamp format; this call was formerly `isc_encode_date()`

These calls merely translate datetime (`DATE`, `TIME`, and `TIMESTAMP`) data between formats; they do not read or write datetime data directly. Datetime data is read from and written to the database using standard DSQL syntax processed with the `isc_dsql` family of API calls.

Note In InterBase 6, the DATE datatype holds only date information in dialect 3 and is not permitted in dialect 1 to avoid ambiguity. When an older database is migrated to version 6 dialect 1, all columns that previously had a DATE datatype are automatically converted to TIMESTAMP. To store migrated data in a DATE column in dialect 3, you must create a new column *in dialect 3* that has the DATE datatype, and then move the data into it. InterBase does not allow you to use ALTER COLUMN to change a TIMESTAMP datatype to a DATE datatype because of potential data loss.

InterBase also requires that numbers entered in database and transaction parameter buffers be in a generic format, with the least significant byte last. Signed numbers require the sign to be in the last byte. Systems that represent numbers with the most significant byte last must use the *isc_vax_integer()* API function to reverse the byte order of numbers entered in database parameter buffers (DPBs) and transaction parameter buffers (TPBs). When numeric information is returned by information calls on these systems, *isc_vax_integer()* must be used once again to reverse the byte ordering.

For more information about using DSQL to read and write data, see **Chapter 6, “Working with Dynamic SQL.”**

Converting date and times from InterBase to C format

The following steps show how to convert the TIMESTAMP datatype from InterBase to C format; the same steps could be used to convert the TIME and DATE datatypes by substituting the appropriate API call above. Starting with InterBase 6, the TIMESTAMP datatype replaces the older DATE datatype used in earlier versions.

To select a timestamp from a table, and convert it to a form usable in a C language program, follow these steps:

1. Create a host variable for a C time structure. Most C and C++ development systems provide a type, *struct tm*, for the C time structure in the **time.h** header file. The following C code includes that header file, and declares a variable of type *struct tm*:

```
#include <time.h>
#include <ibase.h>
. . .
struct tm entry_time;
. . .
```

Note To create host-language time structures in languages other than C and C++, see the host-language reference manual.

2. Create a host variable of type *ISC_TIMESTAMP*. For example, the host-variable declaration might look like this:

```
ISC_TIMESTAMP entry_date;
```

The *ISC_TIMESTAMP* structure is declared in **ibase.h**, but the programmer must declare actual host-language variables of type *ISC_TIMESTAMP*.

3. Retrieve a date from a table into the *ISC_TIMESTAMP* variable.
4. Convert the *ISC_TIMESTAMP* variable into a numeric C format with the InterBase function, *isc_sql_decode_timestamp()*. This function is also declared in **ibase.h**. *isc_sql_decode_timestamp()* requires two parameters, the address of the *ISC_TIMESTAMP* host-language variable, and the address of the *struct tm* host-language variable. For example, the following code fragment converts *entry_date* to *entry_time*:

```
isc_decode_timestamp(&entry_date, &entry_time);
```

Converting dates from C to InterBase format

The following steps show how to convert the *TIMESTAMP* datatype from C to InterBase format; the same steps could be used to convert the *TIME* and *DATE* datatypes by substituting the appropriate API call listed on page 167. To insert a timestamp in a table, it must be converted from the host-language format into InterBase format, and then stored. To perform the conversion and insertion in a C program, follow these steps:

1. Create a host variable for a C time structure. Most C and C++ development systems provide a type, *struct tm*, for the C time structure in the **time.h** header file. The following C code includes that header file, and declares a variable of type *struct tm*:

```
#include <time.h>;
. . .
struct tm entry_time;
. . .
```

To create host-language time structures in languages other than C and C++, see the host-language reference manual.

2. Create a host variable of type *ISC_TIMESTAMP*, for use by InterBase. For example, the host-variable declaration might look like this:

```
ISC_TIMESTAMP mytime;
```

The *ISC_TIMESTAMP* structure is declared in **ibase.h**, but the programmer must declare actual host-language variables of type *ISC_TIMESTAMP*.

3. Put date information into *entry_time*.
4. Use the InterBase *isc_encode_sql_date()* function to convert the information in *entry_time* into InterBase internal format and store that formatted information in the *ISC_TIMESTAMP* host variable (*entry_date* in the example). This function is also declared in **ibase.h**.

isc_encode_sql_timestamp() requires two parameters, the address of the C time structure, and the address of the *ISC_TIMESTAMP* host-language variable. For example, the following code converts *entry_time* to *entry_date*:

```
isc_encode_timestamp(&entry_time, &entry_date);
```

5. Insert the date into a table.

Reversing byte order of numbers with *isc_vax_integer()*

InterBase expects that numbers entered in database and transaction parameter buffers be in a generic format, with the least significant byte last. Signed numbers require the sign to be in the last byte. Systems that represent numbers with the most significant byte last must use the *isc_vax_integer()* API function to reverse the byte order of numbers entered in DPBs and TPBs. When numeric information is returned by information calls on these systems, *isc_vax_integer()* must be used once again to reverse the byte ordering. The syntax for *isc_vax_integer()* is:

```
ISC_LONG isc_vax_integer( char *buffer, short length);
```

buffer is a char pointer to the integer to convert, and *length* is the size, in bytes, of the integer. Valid lengths are 1 (short), 2 (int), and 4(long). The following code reverses the 4-byte value in a result buffer.

```
#include <ibase.h>
. . .
for(p = res_buffer; *p != isc_info_end;)
{
    p++;
    length = isc_vax_integer(p, 2);
}
```

Handling Error Conditions

This chapter describes how to set up an error status vector where InterBase can store run-time error information, and how to use API functions to handle and report errors.

The following table summarizes the API functions for handling errors:

Function	Purpose
<i>isc_interprete()</i>	Capture InterBase error messages to a buffer
<i>isc_print_sqlerror()</i>	Display an SQL error message
<i>isc_print_status()</i>	Display InterBase error messages
<i>isc_sqlcode()</i>	Set the value of SQLCODE
<i>isc_sql_interprete()</i>	Capture an SQL error message to a buffer

TABLE 10.1 Error-handling functions

Setting up an error status vector

Most API functions return status information that indicates success or failure. The information returned is derived from the second array element of the error status vector, where InterBase reports error conditions. The error status vector is declared in applications as an array of 20 long integers, using the following syntax:

```
#include <ibase.h>
. . .
ISC_STATUS status_vector[20];
```

ISC_STATUS is a *#define* in *ibase.h* provided for programing convenience and platform independence.

Using information in the status vector

Whether or not an error occurs during the execution of an API call, InterBase loads the error status vector with status information. Information consists of one or more InterBase error codes, and error information that can be used to build an error message honed to a specific error.

An application can check the status vector after the execution of most API calls to determine their success or failure. If an error condition is reported, applications can:

- Display InterBase error messages using *isc_print_status()*.
- Set an *SQLCODE* value corresponding to an InterBase error using *isc_sqlcode()*, and display the *SQLCODE* and an SQL error message using *isc_print_sqlerror()*.
- Build individual InterBase error messages in a buffer with *isc_interprete()*. The buffer must be provided by the application. Using a buffer enables an application to perform additional message processing (for example, storing messages in an error log file). This ability is especially useful on windowing systems that do not permit direct screen writes.
- Capture an SQL error message in a buffer with *isc_sql_interprete()*. The buffer must be provided by the application.
- Parse for and react to specific InterBase error codes in the status vector.

Checking the status vector for errors

API functions that return information in the status vector are declared in **ibase.h** as returning an *ISC_STATUS* pointer. For example, the function prototype for *isc_prepare_transaction()* is declared as:

```
ISC_STATUS ISC_EXPORT isc_prepare_transaction(
    ISC_STATUS ISC_FAR *,
    isc_tr_handle ISC_FAR *);
```

To check the status vector for error conditions after the execution of a function, examine the first element of the status vector to see if it is set to 1, and if so, examine the second element to see if it is not 0. A nonzero value in the second element indicates an error condition. The following C code fragment illustrates how to check the status vector for an error condition:

```
#include <ibase.h>
. . .
ISC_STATUS status_vector[20];
. . .
/* Assume an API call returning status information is called here. */
if (status_vector[0] == 1 && status_vector[1] > 0)
{
    /* Handle error condition here. */
    ;
}
```

If an error condition is detected, you can use API functions in an error-handling routine to display error messages, capture the error messages in a buffer, or parse the status vector for particular error codes.

Display or capture of error messages is only one part of an error-handling routine. Usually, these routines also roll back transactions, and sometimes they can retry failed operations.

Displaying InterBase error messages

Use *isc_print_status()* to display InterBase error messages on the screen. This function parses the status vector to build all available error messages, then uses the C *printf()* function to write the messages to the display. *isc_print_status()* requires one parameter, a pointer to a status vector containing error information. For example, the following code fragment calls *isc_print_status()* and rolls back a transaction on error:

```
#include <ibase.h>
```

```

. . .
ISC_STATUS status_vector[20];
isc_tr_handle trans;
. . .
trans = 0L;
. . .
/* Assume a transaction, trans, is started here. */
/* Assume an API call returning status information is called here. */
if (status_vector[0] == 1 && status_vector[1] > 0)
{
    isc_print_status(status_vector);
    isc_rollback_transaction(status_vector, &trans);
}

```

IMPORTANT On windowing systems that do not permit direct screen writes with *printf()*, use *isc_interprete()* to capture error messages to a buffer.

Tip For applications that use the dynamic SQL (DSQL) API functions, errors should be displayed using SQL conventions. Use *isc_sqlcode()* and *isc_print_sqlerror()* instead of *isc_print_status()*.

Capturing InterBase error messages

Use *isc_interprete()* to build an error message from information in the status vector and store it in an application-defined buffer where it can be further manipulated. Capturing messages in a buffer is useful when applications:

- Run under windowing systems that do not permit direct screen writes.
- Require more control over message display than is possible with *isc_print_status()*.
- Store a record of all error messages in a log file.
- Manipulate or format error messages for display or pass them to a windowing system's display routines.

isc_interprete() retrieves and formats a single error message each time it is called. When an error occurs, the status vector usually contains more than one error message. To retrieve all relevant error messages, you must make repeated calls to *isc_interprete()*.

Given both the location of a buffer, and the address of the status vector, `isc_interprete()` builds an error message from the information in the status vector, puts the formatted string in the buffer where an application can manipulate it, and advances the status vector pointer to the start of the next cluster of error information. `isc_interprete()` requires two parameters, the address of an application buffer to hold formatted message output, and a pointer to the status vector array.

IMPORTANT Never pass the status vector array directly to `isc_interprete()`. Each time it is called, `isc_interprete()` advances the pointer to the status vector to the next element containing new message information. Before calling `isc_interprete()`, be sure to set the pointer to the starting address of the status vector.

The following code demonstrates an error-handling routine that makes repeated calls to `isc_interprete()` to retrieve error messages from the status vector in a buffer, one at a time, so they can be written to a log file:

```
#include <ibase.h>
. . .
ISC_STATUS status_vector[20];
isc_tr_handle trans;
long *pvector;
char msg[512];
FILE *efile; /* Code fragment assumes pointer to an open file. */
trans = 0L;
. . .
/* Error-handling routine starts here. */
/* Always set pvector to point to start of status_vector. */
pvector = status_vector;
/* Retrieve first message. */
isc_interprete(msg, &pvector);
/* Write first message from buffer to log file. */
fprintf(efile, "%s\n", msg);
msg[0] = '-'; /* Append leading hyphen to secondary messages. */
/* Look for more messages and handle in a loop. */
while(isc_interprete(msg + 1, &pvector)) /* More? */
    fprintf(efile, "%s\n", msg); /* If so, write it to the log. */
fclose(efile); /* All done, so close the log file. */
isc_rollback(status_vector, &trans);
return(1);
. . .
```

Note This code fragment assumes that the log file is properly declared and opened elsewhere in the application *before* control is passed to this error handler.

Tip For applications that use the dynamic SQL (DSQL) API functions, errors should be buffered using SQL conventions. Use *isc_sqlcode()* and *isc_sql_interprete()* instead of *isc_interprete()*.

Setting an SQLCODE value on error

For DSQL applications, error conditions should be cast in terms of SQL conventions. SQL applications typically report errors through a variable, *SQLCODE*, declared by an application. To translate an InterBase error code into *SQLCODE* format, use *isc_sqlcode()*. This function searches the error status vector for an InterBase error code that can be translated into an SQL error code, and performs the translation. Once *SQLCODE* is set, the other API functions for handling SQL errors, *isc_print_sqlerror()*, and *isc_sql_interprete()*, can be called.

isc_sqlcode() requires one parameter, a pointer to the status vector. It returns a long value, containing an SQL error code. The following code illustrates the use of this function:

```
#include <ibase.h>
. . .
long SQLCODE; /* Declare the SQL error code variable. */
ISC_STATUS status_vector[20];
. . .
if (status_vector[0] == 1 && status_vector[1] > 0)
{
    SQLCODE = isc_sqlcode(status_vector);
    isc_print_sqlerror(SQLCODE, status_vector)
    . . .
}
```

If successful, *isc_sqlcode()* returns the first valid SQL error code decoded from the status vector. If no valid SQL error code is found, *isc_sqlcode()* returns -999.

Displaying SQL error messages

API applications that provide a DSQL interface to end users should use *isc_print_sqlerror()* to display SQL error codes and corresponding error messages on the screen. When passed a variable, conventionally named *SQLCODE*, containing an SQL error code, and a pointer to the status vector, *isc_print_sqlerror()* parses the status vector to build an SQL error message, then uses the C *printf()* function to write the *SQLCODE* value and message to the display. For example, the following code fragment calls *isc_print_sqlerror()* and rolls back a transaction on error:

```

#include <ibase.h>
. . .
ISC_STATUS status_vector[20];
isc_tr_handle trans;
long SQLCODE;
. . .
trans = 0L;
. . .
/* Assume a transaction, trans, is started here. */
/* Assume an API call returning status information is called here. */
if (status_vector[0] == 1 && status_vector[1] > 0)
{
    SQLCODE = isc_sqlcode(status_vector);
    isc_print_sqlerror(SQLCODE, status_vector);
    isc_rollback_transaction(status_vector, &trans);
}

```

IMPORTANT On windowing systems that do not permit direct screen writes with *printf()*, use *isc_sql_interprete()* to capture error messages to a buffer.

Capturing SQL error messages

Use *isc_sql_interprete()* to build an SQL error message based on a specific SQL error code and store it in a buffer defined by an application. Capturing messages in a buffer is useful when applications:

- Run under windowing systems that do not permit direct screen writes.
- Store a record of all error messages in a log file.
- Manipulate or format error messages for display or pass them to a windowing system's display routines.

isc_sql_interprete() requires three parameters: a valid SQL error code, usually passed as a variable named *SQLCODE*, a buffer where the SQL message should be stored, and the size of the buffer. The following code illustrates how this function might be called to build a message string and store it in a log file:

```

#include <ibase.h>
. . .
ISC_STATUS status_vector[20];
isc_tr_handle trans;
long SQLCODE;
char msg[512];

```

```

FILE *efile; /* Code fragment assumes pointer to an open file. */
trans = 0L;

. . .
/* Assume a transaction, trans, is started here. */
/* Assume an API call returning status information is called here. */
. . .
/* Error-handling routine starts here. */
if (status_vector[0] == 1 && status_vector[1] > 0)
{
    SQLCODE = isc_sqlcode(status_vector);
    isc_sql_interprete(SQLCODE, msg, 512);
    fprintf(efile, "%s\n", msg);
    isc_rollback_transaction(status_vector, &trans);
    return(1);
}

```

Note This code fragment assumes that the log file is properly declared and opened elsewhere in the application *before* control is passed to this error handler.

Parsing the status vector

InterBase stores error information in the status vector in clusters of two or three longs. The first cluster in the status vector *always* indicates the primary cause of the error. Subsequent clusters may contain supporting information about the error, for example, strings or numbers for display in an associated error message. The actual number of clusters used to report supporting information varies from error to error.

In many cases, additional errors may be reported in the status vector. Additional errors are reported immediately following the first error and its supporting information, if any. The first cluster for each additional error message identifies the error. Subsequent clusters may contain supporting information about the error.

► *How the status vector is parsed*

The InterBase error-handling routines, `isc_print_status()` and `isc_interprete()`, use routines which automatically parse error message information in the status vector without requiring you to know about its structure. If you plan to write your own routines to read and react to the contents of the status vector, you need to know how to interpret it.

The key to parsing the status vector is to decipher the meaning of the first long in each cluster, beginning with the first cluster in the vector.

► *Meaning of the first long in a cluster*

The first long in any cluster is a *numeric descriptor*. By examining the numeric descriptor for any cluster, you can always determine the:

- Total number of longs in the cluster.
- Kind of information reported in the remainder of the cluster.
- Starting location of the next cluster in the status vector.

Interpretation of 1st long in a cluster		
Value	Longs in cluster	Meaning
0	—	End of error information in the status vector
1	2	Second long is an InterBase error code
2	2	Second long is the address of string used as a replaceable parameter in a generic InterBase error message
3	3	Second long is the length, in bytes, of a variable-length string provided by the operating system (most often this string is a file name); third long is the address of the string
4	2	Second long is a number used as a replaceable parameter in a generic InterBase error message
5	2	Second long is the address of an error message string requiring no further processing before display
6	2	Second long is a VAX/VMS error code
7	2	Second long is a UNIX error code
8	2	Second long is an Apollo Domain error code

TABLE 10.2 Interpretation of status vector clusters

Interpretation of 1st long in a cluster

Value	Longs in cluster	Meaning
9	2	Second long is an MS-DOS or OS/2 error code.
10	2	Second long is an HP MPE/XL error code.
11	2	Second long is an HP MPE/XL IPC error code.
12	2	Second long is a NeXT/Mach error code.

Note: As InterBase is adapted to run on other hardware and software platforms, additional numeric descriptors for specific platform and operating system error codes will be added to the end of this list.

TABLE 10.2 Interpretation of status vector clusters (*continued*)

By including **ibase.h** at the start of your source code, you can use a series of *#defines* to substitute for hard-coded numeric descriptors in the status vector parsing routines you write. The advantages of using these *#defines* over hard-coding the descriptors are:

- Your code will be easier to read.
- Code maintenance will be easier should the numbering scheme for numeric descriptors change in a future release of InterBase.

The following table lists the *#define* equivalents of each numeric descriptor:

Value	#define	Value	#define
0	<i>isc_arg_end</i>	8	<i>isc_arg_domain</i>
1	<i>isc_arg_gds</i>	9	<i>isc_arg_dos</i>
2	<i>isc_arg_string</i>	10	<i>isc_arg_mpexl</i>
3	<i>isc_arg_cstring</i>	11	<i>isc_arg_mpexl_ipc</i>
4	<i>isc_arg_number</i>	15	<i>isc_arg_next_mach</i>
5	<i>isc_arg_interpreted</i>	16	<i>isc_arg_netware</i>
6	<i>isc_arg_vms</i>	17	<i>isc_arg_win32</i>
7	<i>isc_arg_unix</i>		

TABLE 10.3 *#defines* for status vector numeric descriptors

For an example of code that uses these defines, see **“Status vector parsing example” on page 183**.

► *Meaning of the second long in a cluster*

The second long in a cluster is *always* one of five items:

- An InterBase error code (1st long = 1).
- A string address (1st long = 2 or 5).
- A string length (1st long = 3).
- A numeric value (1st long = 4).
- An operating system error code (1st long > 5).

INTERBASE ERROR CODES

InterBase error codes have two uses. First, they are used internally by InterBase functions to build and display descriptive error message strings. For example, *isc_interprete()* calls another function which uses the InterBase error code to retrieve a base error message from which it builds an error message string you can display or store in a log file.

Secondly, when you write your own error-handling routine, you can examine the status vector directly, trapping for and reacting to specific InterBase error codes.

When the second long of a cluster is an InterBase error code, then subsequent clusters may contain additional parameters for the error message string associated with the error code. For example, a generic InterBase error message may contain a replaceable string parameter for the name of the table where an error occurs, or it may contain a replaceable numeric parameter for the code of the trigger which trapped the error condition.

If you write your own parsing routines, you may need to examine and use these additional clusters of error information.

STRING ADDRESSES

String addresses point to error message text. When the first long in the cluster is 2 (*isc_arg_string*), the address pointed to often contains the name of the database, table, or column affected by the error. In these cases, InterBase functions which build error message strings replace a parameter in a generic InterBase error message with the string pointed to by this address. Other times the address points to an error message hard-coded in a database trigger.

When the first long in the cluster is 5 (*isc_arg_interpreted*), the address points to a text message which requires no further processing before retrieval. Sometimes this message may be hard-coded in InterBase itself, and other times it may be a system-level error message.

In either of these cases, InterBase functions such as *isc_print_status()* and *isc_interprete()* can format and display the resulting error message for you.

STRING LENGTH INDICATORS

When the first long in a cluster is 3 (*isc_arg_cstring*), the numeric value in the second long indicates the length, in bytes, of a message string whose address is stored in the third long in the cluster. This string requires translation into a standard, null-terminated C string before display.

NUMERIC VALUES

A numeric value has different meaning depending upon the value of the numeric descriptor in the first long of a cluster. If the first long is 4 (*isc_arg_number*), a numeric value is used by InterBase functions to replace numeric parameters in generic InterBase error messages during message building. For instance, when an integrity error occurs, InterBase stores the code of the trigger which detects the problem as a numeric value in the status vector. When an InterBase function like *isc_interprete()* builds the error message string for this error, it inserts the numeric value from the status vector into the generic InterBase integrity error message string to make it more specific.

OPERATING SYSTEM ERROR CODES

If the first long in a cluster is greater than 5, the numeric value in the second long is an error code specific to a particular platform or operating system. InterBase functions should not be used to retrieve and display the specific platform or operating system error message. Consult your operating system manual for information on how to handle such errors.

► *Meaning of the third long in a cluster*

If the first long in a cluster is 3 (*isc_arg_cstring*), the cluster's total length is three longs. The third long *always* contains the address of a message string requiring translation into a standard, null-terminated C string before display. Such a string is often a file or path name. InterBase functions like *isc_interprete()* automatically handle this translation for you.

► *Status vector parsing example*

The following C example illustrates a simple, brute force parsing of the status vector. The code forces an error condition. The error-handling block parses the status vector array cluster by cluster, printing the contents of each cluster and interpreting it for you.

```
#include <ibase.h>
. . .
ISC_STATUS status_vector[20];
main()
{
    int done, v; /* end of args?, index into vector */
    int c, extra; /* cluster count, 3 long cluster flag */
    static char *meaning[] = {"End of error information",
        "n InterBase error code", " string address", " string length",
        " numeric value", " system code"};
    /* Assume database is connected and transaction started here. */
    if (status_vector[0] == 1 && status_vector[1] > 0)
        error_exit();
    . . .
}
void error_exit(void)
{
    done = v = 0;
    c = 1;
    while (!done)
    {
        extra = 0;
        printf("Cluster %d:\n", c);
        printf("Status vector %d: %ld: ", v, status_vector[v]);
        if (status_vector[v] != gds_arg_end)
            printf("Next long is a");
        switch (status_vector[v++])
        {
            case gds_arg_end:
                printf("%s\n", meaning[0]);
                done = 1;
                break;
            case gds_arg_gds:
                printf("%s\n", meaning[1]);
                break;
            case gds_arg_string:
            case gds_arg_interpreted:
```

```

        printf("%s\n", meaning[2]);
        break;
    case gds_arg_number:
        printf("%s\n", meaning[4]);
        break;
    case gds_arg_cstring:
        printf("%s\n", meaning[3]);
        extra = 1;
        break;
    default:
        printf("%s\n", meaning[5]);
        break;
}
if (!done)
{
    printf("Status vector %d: %ld", v, status_vector[v]);
    v++;/* advance vector pointer */
    c++;/* advance cluster count */
    if (extra)
    {
        printf(": Next long is a %s\n", meaning[2]);
        printf("Status vector: %d: %ld\n\n", v,
            status_vector[v]);
        v++;
    }
    else
        printf("\n\n");
}
}
isc_rollback_transaction(status_vector, &trans);
isc_detach_database(&dbl);
return(1);
}

```

Here is a sample of the output from this program:

```

Cluster 1:
Status vector 0: 1: Next long is an InterBase error code
Status vector 1: 335544342
Cluster 2:
Status vector 2: 4: Next long is a numeric value
Status vector 3: 1
Cluster 3:

```

```

Status vector 4: 1: Next long is an InterBase error code
Status vector 5: 335544382
Cluster 4:
Status vector 6: 2: Next long is a string address
Status vector 7: 156740
Cluster 5:
Status vector 8: 0: End of error information

```

This output indicates that two InterBase errors occurred. The first error code is 335544342. The error printing routines, *isc_print_status()* and *isc_interprete()*, use the InterBase error code to retrieve a corresponding base error message. The base error message contains placeholders for replaceable parameters. For error code 335544342, the base error message string is:

```
"action cancelled by trigger (%ld) to preserve data integrity"
```

This error message uses a replaceable numeric parameter, *%ld*.

In this case, the numeric value to use for replacement, 1, is stored in the second long of the second cluster. When the error printing routine inserts the parameter into the message, it displays the message:

```
action cancelled by trigger (1) to preserve data integrity
```

The second error code is 335544382. The base message retrieved for this error code is:

```
"%s"
```

In this case, the entire message to be displayed consists of a replaceable string. The second long of the fourth cluster contains the address of the replacement string, 156740. This is an error message defined in the trigger that caused the error. When the error printing routine inserts the message from the trigger into the base message, it displays the resulting message:

```
-Department name is missing.
```

Note This sample program is only meant to illustrate the structure of the status vector and its contents. While the error-handling routine in this program might serve as a limited debugging tool for a program under development, it does not provide useful information for end users. Ordinarily, error-handling blocks in applications should interpret errors, display explanatory error messages, and take corrective action, if appropriate.

For example, if the error-handling routine in the sample program had called *isc_print_status()* to display the error messages associated with these codes, the following messages would have been displayed:

```
action cancelled by trigger (1) to preserve data integrity
-Department name is missing.
```


11

Working with Events

This chapter describes how to work with *events*, a message passed from a trigger or stored procedure to an application to announce the occurrence of a specified condition or action, usually a database change such as an insertion, modification, or deletion of a record. It explains how to set up event buffers, and use the following API functions to make synchronous and asynchronous event calls. In the following table, functions are listed in the order they typically appear in an application:

Function	Purpose
<i>isc_event_block()</i>	Allocate event parameter buffers
<i>isc_wait_for_event()</i>	Wait for a synchronous event to be posted
<i>isc_que_events()</i>	Set up an asynchronous event and return to application processing
<i>isc_event_counts()</i>	Determine the change in values of event counters in the event parameter buffer
<i>isc_cancel_events()</i>	Cancel interest in an event

TABLE 11.1 API event functions

For asynchronous events, this chapter also describes how to create an asynchronous trap (AST), a function that responds to posted events.

Understanding the event mechanism

The InterBase event mechanism consists of four parts:

- The InterBase engine that maintains an event queue and notifies applications when an event occurs.
- Event parameter buffers set up by an application where it can receive notification of events.
- An application that registers interest in one or more specified, named events and either waits for notification to occur (synchronous event), or passes a pointer to an AST function that handles notifications so that application processing can continue in the meantime (asynchronous event).
- A trigger or stored procedure that notifies the engine that a specific, named event has occurred. Notification is called *posting*.

The InterBase event mechanism enables applications to respond to actions and database changes made by other, concurrently running applications without the need for those applications to communicate directly with one another, and without incurring the expense of CPU time required for periodic polling to determine if an event has occurred.

For information about creating triggers and stored procedures that post events, see the *Data Definition Guide*.

Event parameter buffers

If an application is to receive notification about events, it must set up two identically-sized event parameter buffers (EPBs) using `isc_event_block()`. The first buffer, *event_buffer*, is used to hold the count of event occurrences before the application registers an interest in the event. The second buffer, *result_buffer*, is subsequently filled in with an updated count of event occurrences when an event of interest to the application occurs. A second API function, `isc_event_counts()`, determines the differences between item counts in these buffers to determine which event or events occurred.

For more information about setting up and using EPBs, see **“Creating EPBs with `isc_event_block()`” on page 190**.

Synchronous event notification

When an application depends on the occurrence of a specific event for processing, it should use synchronous event notification to suspend its own execution until the event occurs. For example, an automated stock trading application that buys or sells stock when specific price changes occur might start execution, set up EPBs, register interest in a set of stocks, then suspend its own execution until those price changes occur.

The `isc_wait_for_event()` function provides synchronous event handling for an application. For more information about synchronous event handling, see **“Waiting on events with `isc_wait_for_event()`” on page 191**.

Asynchronous event notification

When an application needs to react to possible database events, but also needs to continue processing whether or not those events occur, it should set up an asynchronous trap (AST) function, and use asynchronous event notification to register interest in events while continuing its own processing. For example, a stock brokering application requires constant access to a database of stocks to allow a broker to buy and sell stock, but, at the same time, may want to use events to alert the broker to particularly significant or volatile stock price changes.

The `isc_que_events()` function and the AST function provide asynchronous event handling for an application. For more information about asynchronous event handling, see **“Continuous processing with `isc_que_events()`” on page 192**.

Transaction control of events

Events occur under transaction control, and can therefore be committed or rolled back. Interested applications do not receive notification of an event until the transaction from which the event is posted is committed. If a posted event is rolled back, notification does not occur.

A transaction can post the same event more than once before committing, but regardless of how many times an event is posted, it is regarded as a single event occurrence for purposes of event notification.

Creating EPBs with *isc_event_block()*

Before an application can register interest in an event, it must establish and populate two event parameter buffers (EPBs), one for holding the initial occurrence count values for each event of interest, and another for holding the changed occurrence count values. These buffers are passed as parameters to several API event functions.

In C, each EPB is declared as a char pointer, as follows:

```
char *event_buffer, *result_buffer;
```

Once the buffers are declared, *isc_event_block()* is called to allocate space for them, and to populate them with starting values.

isc_event_block() also requires at least two additional parameters: the number of events in which an application is registering interest, and, for each event, a string naming the event. A single call to *isc_event_block()* can pass up to 15 event name strings. Event names must match the names of events posted by stored procedures or triggers.

isc_event_block() allocates the same amount of space for each EPB, enough to handle each named event. It returns a single value, indicating the size, in bytes, of each buffer.

The syntax for *isc_event_block()* is:

```
ISC_LONG isc_event_block(
char **event_buffer,
char **result_buffer,
unsigned short id_count,
. . . );
```

For example, the following code sets up EPBs for three events:

```
#include <ibase.h>
. . .
char *event_buffer, *result_buffer;
long blength;
. . .
blength = isc_event_block(&event_buffer, &result_buffer, 3, "BORL",
"INTEL", "SUN");
. . .
```

This code assumes that there are triggers or stored procedures defined for the database that post events named “BORL”, “INTEL”, and “SUN”.

Tip Applications that need to respond to more than 15 events can make multiple calls to *isc_event_block()*, specifying different EPBs and event lists for each call.

For the complete syntax of *isc_event_block()*, see **“isc_event_block()” on page 347**.

Waiting on events with `isc_wait_for_event()`

After setting up EPBs and specifying events of interest with `isc_event_block()`, an application can use `isc_wait_for_event()` to register interest in those events and pause its execution until one of the events occurs.

IMPORTANT `isc_wait_for_event()` cannot be used in Microsoft Windows applications or under any other operating system that does not permit processes to pause. Applications on these platforms must use asynchronous event handling.

The syntax for `isc_wait_for_event()` is:

```
ISC_STATUS isc_wait_for_event(
    ISC_STATUS *status_vector,
    isc_db_handle *db_handle,
    short length,
    char *event_buffer,
    char *result_buffer);
```

For example, the following code sets up EPBs for three events, then calls `isc_wait_for_event()` to suspend its execution until one of the events occurs:

```
#include <ibase.h>
. . .
char *event_buffer, *result_buffer;
long length;
ISC_STATUS status_vector[20];
isc_db_handle db1;
. . .
/* Assume database db1 is attached here and a transaction started. */
length = isc_event_block(&event_buffer, &result_buffer, 3, "BORL",
    "INTEL", "SUN");
isc_wait_for_event(status_vector, &db1, (short)length,
    event_buffer, result_buffer);
/* Application processing is suspended here until an event occurs. */
. . .
```

Once `isc_wait_for_event()` is called, application processing stops until one of the requested events is posted. When the event occurs, application processing resumes at the next executable statement following the call to `isc_wait_for_event()`. If an application is waiting on more than one event, it must use `isc_event_counts()` to determine which event was posted.

Note A single call to `isc_wait_for_event()` can wait on a maximum of 15 events. Applications that need to wait on more than 15 events must wait on one set of 15, then make another call to `isc_wait_for_event()` to wait on additional events.

For the complete syntax of `isc_wait_for_event()`, see **“`isc_wait_for_event()`” on page 394**.

Continuous processing with `isc_que_events()`

`isc_que_events()` is called to request asynchronous notification of events listed in an event buffer passed as an argument. Upon completion of the call, but before any events are posted, control is returned to the calling application so that it can continue processing.

When a requested event is posted, InterBase calls an asynchronous trap (AST) function, also passed as a parameter to `isc_que_events()`, to handle the posting. The AST is a function or subroutine in the calling application, the sole purpose of which is to process the event posting for the application.

The syntax for `isc_que_events()` is:

```
ISC_STATUS isc_que_events(
    ISC_STATUS *status_vector,
    isc_db_handle *db_handle,
    ISC_LONG *event_id,
    short length,
    char *event_buffer,
    isc_callback event_function,
    void *event_function_arg);
```

`event_id` is a long pointer that is used as a handle in subsequent calls to `isc_cancel_events()` to terminate event notification. It need not be initialized when passed. The `length` parameter is the size of `event_buffer`, which contains the current count of events to be waited upon. `event_function` is a pointer to the AST function that InterBase should call when an event of interest is posted. It is up to the AST function to notify the application that it has been called, perhaps by setting a global flag of some kind. `event_function_arg` is a pointer to the first parameter to pass to the AST.

For a complete example of a call to `isc_que_events()` and a call to an AST, see **“A complete `isc_que_events()` example” on page 193**.

Creating an AST

The event function, *event_function*, should be written to take three arguments:

1. The *event_function_arg* specified in the call to `isc_que_events()`. This is usually a pointer to the event parameter buffer that should be filled in with updated event counts.
2. The length of the following *events_list* buffer.
3. A pointer to the *events_list* buffer, a temporary event parameter buffer just like that passed to `isc_que_events()`, except for having updated event counts.

A result buffer is not automatically updated by the event occurrence; it is up to the *event_function* to copy the temporary *events_list* buffer to the more permanent buffer that the application utilizes.

event_function also needs to let the application know that it has been called, for example, by setting a global flag.

A sample *event_function* appears in the following example:

```
isc_callback event_function
    (char *result, short length, char *updated)
{
    /* Set the global event flag. */
    event_flag++
    /* Copy the temporary updated buffer to the result buffer. */
    while (length-- > 0)
        *result++ = *updated++;
    return(0);
};
```

A complete `isc_que_events()` example

The following program fragment illustrates calling `isc_que_events()` to wait asynchronously for event occurrences. Within a loop, it performs other processing, and checks the event flag (presumably set by the specified event function) to determine when an event has been posted. If one has, the program resets the event flag, calls `isc_event_counts()` to determine which events have been posted since the last call to `isc_que_events()`, and calls `isc_que_events()` to initiate another asynchronous wait.

```
#include <ibase.h>
#define number_of_stocks 3;
#define MAX_LOOP 10
char *event_names[] = {"DEC", "HP", "SUN"};
```

```

char *event_buffer, *result_buffer;
ISC_STATUS status_vector[20];
short length;
ISC_LONG event_id;
int i, counter;
int event_flag = 0;
length = (short)isc_event_block(
    &event_buffer,
    &result_buffer,
    number_of_stocks,
    "DEC", "HP", "SUN");
isc_que_events(
    status_vector,
    &database_handle, /* Set in previous isc_attach_database(). */
    &event_id,
    length, /* Returned from isc_event_block(). */
    event_buffer,
    (isc_callback)event_function,
    result_buffer);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector); /* Display error message. */
    return(1);
};
counter = 0;
while (counter < MAX_LOOP)
{
    counter++;
    if (!event_flag)
    {
        /* Do whatever other processing you want. */
        ;
    }
    else
    {
        event_flag = 0;
        isc_event_counts(
            status_vector,
            length,
            event_buffer,
            result_buffer);
        if (status_vector[0] == 1 && status_vector[1])
        {

```

```

        isc_print_status(status_vector); /*Display error message.*/
        return(1);
    };
    for (i=0; i<number_of_stocks; i++)
        if (status_vector[i])
        {
            /* The event has been posted. Do whatever is appropriate,
               such as initiating a buy or sell order.
               Note: event_names[i] tells the name of the event
               corresponding to status_vector[i]. */

            ;
        }
    isc_que_events(
        status_vector,
        &database_handle,
        &event_id,
        length,
        event_buffer,
        (isc_callback)event_function,
        result_buffer);
    if (status_vector[0] == 1 && status_vector[1])
    {
        isc_print_status(status_vector); /*Display error message.*/
        return(1);
    }
} /* End of else. */
} /* End of while. */
/* Let InterBase know you no longer want to wait asynchronously. */
isc_cancel_events(
    status_vector,
    &database_handle,
    &event_id);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector); /* Display error message. */
    return(1);
}

```

Determining which events occurred with *isc_event_counts()*

When an application registers interest in multiple events and receives notification that an event occurred, the application must use *isc_event_counts()* to determine which event or events occurred. *isc_event_counts()* subtracts values in the *event_buffer* array from the values in the *result_buffer* array to determine the number of times each event has occurred since an application registered interest in a set of events. *event_buffer* and *result_buffer* are variables declared within an application, and allocated and initialized by *isc_event_block()*.

The difference of each element is returned in the error status array that is passed to *isc_event_counts()*. To determine which events occurred, an application must examine each element of the array for nonzero values. A nonzero count indicates the number of times an event is posted between the time *isc_event_block()* is called and the first time an event is posted after *isc_wait_for_event()* or *isc_que_events()* are called. Where multiple applications are accessing the same database, therefore, a particular event count may be 1 or more, and more than one event count element may be nonzero.

Note When first setting up an AST to trap events with *isc_que_events()*, InterBase initializes all count values in the status vector to 1, rather than 0. To clear the values, call *isc_event_counts()* to reset the values.

In addition to determining which event occurred, *isc_event_counts()* reinitializes the *event_buffer* array in anticipation of another call to *isc_wait_for_event()* or *isc_que_events()*. Values in *event_buffer* are set to the same values as corresponding values in *result_buffer*.

The complete syntax for *isc_event_counts()* is:

```
void isc_event_counts(
    ISC_STATUS status_vector,
    short buffer_length,
    char *event_buffer,
    char *result_buffer);
```

For example, the following code declares interest in three events, waits on them, then uses *isc_event_counts()* to determine which events occurred:

```
#include <ibase.h>;
. . .
char *event_buffer, *result_buffer;
long blength;
ISC_STATUS status_vector[20];
isc_db_handle db1;
long count_array[3];
int i;
```

```

. . .
/* Assume database db1 is attached here and a transaction started. */
blength = isc_event_block(&event_buffer, &result_buffer, 3, "BORL",
    "INTEL", "SUN");
isc_wait_for_event(status_vector, &db1, (short)blength,
    event_buffer, result_buffer);
/* Application processing is suspended here until an event occurs. */
isc_event_counts(status_vector, (short)blength, event_buffer,
    result_buffer);
for (i = 0; i < 3; i++)
{
    if (status_vector[i])
    {
        /* Process the event here. */
        ;
    }
}

```

For more information about `isc_event_counts()`, see **“`isc_event_counts()`” on page 349** of **Chapter 13, “API Function Reference.”**

Canceling interest in asynchronous events with `isc_cancel_events()`

An application that requested asynchronous event notification with `isc_que_events()` can subsequently cancel the notification request at any time with `isc_cancel_events()` using the following syntax:

```

ISC_STATUS isc_cancel_events(
    ISC_STATUS *status_vector,
    isc_db_handle *db_handle,
    ISC_LONG *event_id);

```

event_id is an event handle set in a previous call to `isc_que_events()`. For example, the following code cancels interest in the event or events identified by *event_id*:

```

include <ibase.h>;

. . .
/* For example code leading up to this call, see the code example
in "Continuous Processing with isc_event_que()", earlier in this
chapter. */
isc_cancel_events(status_vector, &db_handle, &event_id);

```



12

Working with Services

This chapter covers the InterBase Services API functions. This facility allows you to write applications that monitor and control InterBase servers and databases. Tasks that you can perform with this API include:

- Performing database maintenance tasks such as database backup and restore, shutdown and restart, garbage collection, and scanning for invalid data structures
- Creating, modifying, and removing user entries in the security database
- Administering software activation certificates
- Requesting information about the configuration of databases and the server

Overview of the Services API

This section describes general concepts of the Services API, usage of the services parameter buffer, and methods for attaching and detaching from a Services Manager.

General information

The Services API is a group of functions in the InterBase client library (**gds32.dll** on Windows, **libgds.a** on UNIX/Linux). The features that you can exercise with the Services API include those of the command-line tools **gbak**, **gfix**, **gsec**, **gstat**, and **iblicense** (see the *Operations Guide* for information on these tools). The Services API can also perform other functions that are not provided by these tools.

All InterBase servers include a facility called the *Services Manager*. The Services API enables client applications to submit requests to the Services Manager of an InterBase server, and the Services Manager performs the tasks. The server can be local (on the same host as your application), or remote (on another host on the network). The Services API offers the same features when connected to either local or remote InterBase servers.

The Services API family consists of the following four functions:

- *isc_service_attach()* initiates a connection to a specified Services Manager
- *isc_service_start()* invokes a service task
- *isc_service_query()* requests information or task results from the Services Manager
- *isc_service_detach()* disconnects from the Services Manager

For full details on the syntax and options of the Services API functions, see the reference entries for **“isc_service_attach()” on page 376**, **“isc_service_detach()” on page 377**, **“isc_service_query()” on page 378**, and **“isc_service_start()” on page 380**.

Using services parameter buffers

You can specify options to tailor your attachment to a Services Manager by creating a services parameter buffer (SPB), populating it with desired properties, and passing the address of the SPB to *isc_service_attach()* or other functions in the Services API group. For example, the SPB can contain a user name and password for attaching to a remote server.

An SPB is a *char* array variable that you declare in your application. It contains the following elements:

1. A byte that introduces the version of the SPB format, always the compile-time constant, *isc_spb_version*.
2. A byte that specifies the version number. InterBase supplies a macro *isc_spb_current_version*, that is defined as the recommended SPB version for each given release of the InterBase product.
3. A contiguous series of one or more *clusters* of bytes follow, each describing a single argument.

Each cluster consists of the following parts:

1. A byte that introduces the parameter type for each cluster. There are compile-time constants defined for all the parameter types (for example, *isc_spb_user_name*).
2. A byte that specifies the number of bytes that follow in the remainder of the cluster argument; this is not needed for certain parameter types that have fixed-length arguments.
3. A variable number of bytes that contain data, depending on the parameter type.

Subsequent clusters follow immediately in the SPB array.

For example, the following C/C++ code fills an SPB buffer with the SPB version and a cluster for the user name.

EXAMPLE 12.1 Filling a services parameter buffer in C/C++

```

1 char spb_buffer[128], *spb = spb_buffer;
2 *spb++ = isc_spb_version;
3 *spb++ = isc_spb_current_version;
4 *spb++ = isc_spb_user_name;
5 *spb++ = strlen("SYSDBA");
6 strcpy(spb, "SYSDBA");
7 spb += strlen("SYSDBA");

```

Line 1 declares an array of 128 bytes, and a pointer initialized to the first entry in the array.

Line 2 assigns the item specifier for the SPB version to the first element of the array. Every SPB must have this item at the start of the array. Since this SPB item is always one byte long, it doesn't take a length specifier.

Line 3 assigns the value for the SPB version item.

Line 4 assigns the cluster identifier for the user name string to the next element of the array.

Line 5 provides the length of the following string. In this example, the string is “SYSDBA”, and the length is 6.

Line 6 copies the string “SYSDBA” into the array starting at the current element.

Line 7 increments the SPB pointer past the string “SYSDBA”, positioning it for additional clusters.

IMPORTANT All numbers in the database parameter buffer must be represented in a generic format, with the least significant byte first. Signed numbers should have the sign in the last byte. The API function `isc_vax_integer()` can be used to reverse the byte order of a number. For more information, see **“isc_vax_integer()” on page 391**.

Attaching to the Services Manager with `isc_service_attach()`

Use the Services API function `isc_service_attach()` to initiate a connection from your application to a remote InterBase Services Manager.

You must supply a local or remote *service name* to specify which host to attach. This string resembles InterBase database connection strings, in that the syntax determines the network protocol used to connect the client application to the Services Manager on the server host.

Protocol	Syntax	Supported server platform
TCP/IP	<code>serverhost:service_mgr</code>	any
NetBEUI	<code>\\serverhost\service_mgr</code>	Windows NT
IPX/SPX	<code>serverhost@service_mgr</code>	NetWare
Local	<code>service_mgr</code>	any

TABLE 12.1 Syntax of Services Manager connect string, by protocol

Replace *serverhost* with the hostname of your InterBase database server. In all cases, the string *service_mgr* is a literal string.

The user ID you use to attach to the Services Manager is the user ID the Services Manager uses to perform tasks by your request. Note that some service tasks can be performed only by the SYSDBA user ID.

EXAMPLE 12.2 Attaching to a Services Manager in C/C++

```

char *user = "SYSDBA",
    *password = "masterkey", /* see security tip below */
    *service_name = "jupiter:service_mgr";
ISC_STATUS status[20];
isc_svc_handle *service_handle = NULL;
spb_buffer[128], *spb = spb_buffer;
unsigned short spb_length;

*spb++ = isc_spb_version;
*spb++ = isc_spb_current_version;

*spb++ = isc_spb_user_name;
*spb++ = strlen(user);
strcpy(spb, user);
spb += strlen(user);

*spb++ = isc_spb_password;
*spb++ = strlen(password);
strcpy(spb, password);
spb += strlen(password);

spb_length = spb - spb_buffer;

if (isc_service_attach(status, 0, service_name,
    &service_handle, spb_length, spb_buffer))
{
    isc_print_status(status);
    exit(1);
}

```

Detaching from a Services Manager with *isc_service_detach()*

Use *isc_service_detach()* after you finish your tasks with the Services API, to end the connection with the Services Manager. Following is a C/C++ code example of terminating the connection, assuming you have acquired a valid service handle from *isc_service_attach()*.

EXAMPLE 12.3 Detaching from a Services Manager in C/C++

```
isc_service_detach(status, &service_handle);
```

Invoking service tasks with *isc_service_start()*

You can use the function *isc_service_start()* to request that the Services Manager perform specified tasks. These tasks execute on the server host as a thread in the **ibserver** process. This section describes the types of tasks you can request.

You can execute only one task at a time in a given attachment to a Services Manager. While the task is running, you can retrieve any output of the task using *isc_service_query()*. You can maintain multiple attachments to a Services Manager and execute a task in each attachment.

Using request buffers

The Services API uses a buffer structured similarly to the SPB for *isc_service_start()* to specify tasks and options for the Services Manager. This is called the *request buffer*. You supply clusters of parameters and arguments in the request buffer. The Services Manager performs tasks you specify.

Overview of task identifiers

The following table lists by request buffer cluster identifier the tasks that you can request with *isc_service_start()*.

Task item	Purpose
<i>isc_action_svc_backup</i>	Back up a database to a file or tape device; equivalent to gbak -b
<i>isc_action_svc_restore</i>	Restore a database backup file and recreate a database; equivalent to gbak -c
<i>isc_action_svc_properties</i>	Set database properties; equivalent to gfix with various options
<i>isc_action_svc_repair</i>	Initiate database consistency check and correction; equivalent to gfix with -validate , -full , and -mend options
<i>isc_action_svc_db_stats</i>	Report database statistics; equivalent to the output of gstat
<i>isc_action_svc_get_ib_log</i>	Report contents of the interbase.log file on the server
<i>isc_action_svc_display_users</i>	Display a user entry to the security database on the server; equivalent to gsec -display

TABLE 12.2 Services API tasks

Task item	Purpose
<code>isc_action_svc_add_user</code>	Add a user entry to the security database on the server; equivalent to <code>gsec -add</code>
<code>isc_action_svc_delete_user</code>	Delete a user entry to the security database on the server; equivalent to <code>gsec -delete</code>
<code>isc_action_svc_modify_user</code>	Modify a user entry to the security database on the server; equivalent to <code>gsec -modify</code>
<code>isc_action_svc_add_license</code>	Add a software activation certificate to <code>ib_license.dat</code> ; only SYSDBA can invoke this task
<code>isc_action_svc_remove_license</code>	Remove a software activation certificate from <code>ib_license.dat</code> ; only SYSDBA can invoke this task

TABLE 12.2 Services API tasks

See the following sections for descriptions of tasks and examples of starting them.

Backing up and restoring databases

Use the cluster identifier `isc_action_svc_backup` to request that the Services Manager perform a backup operation. This is a programmatic method to invoke the `gbak` tool as a thread in the `ibserver` process. You must specify the path of the database primary file, and the path of the backup output file or files.

Note Paths of backup files are relative to the server. Since the Services Manager executes backup and restore tasks on the server host, the Services Manager reads and writes backup files on the server host.

You can specify additional options as needed. Some options require arguments, while other options are bits in an option bitmask.

The following table lists arguments to `isc_action_svc_backup`:

Argument	Purpose	Argument length	Argument value
<i>isc_spb_dbname</i>	Path of the primary file of the database, from the server's point of view	2-byte length + string	String
<i>isc_spb_verbose</i>	If specified, the Services Manager prepares output to return via <i>isc_service_query()</i> ; corresponds to gbak -verbose	—	—
<i>isc_spb_bkp_file</i>	Path of a backup output file; you can specify multiple output files; corresponds to gsplit functionality	2-byte length + string	String
<i>isc_spb_bkp_length</i>	Length in bytes of the backup output file; you must specify one length value for each output file except the last; corresponds to gsplit functionality	2-byte length + string	String
<i>isc_spb_bkp_factor</i>	Tape device blocking factor; corresponds to gbak -factor	4 bytes	Unsigned long
<i>isc_spb_options</i>	The following value is a bitmask of <i>isc_spb_bkp_xxxx</i> options below	4 bytes	Bitmask
<i>isc_spb_bkp_ignore_checksums</i>	Ignore checksums during backup; corresponds to gbak -ignore	—	Bit
<i>isc_spb_bkp_ignore_limbo</i>	Ignore limbo transactions during backup; corresponds to gbak -limbo	—	Bit
<i>isc_spb_bkp_metadata_only</i>	Output backup file for metadata only with empty tables; corresponds to gbak -metadata	—	Bit
<i>isc_spb_bkp_no_garbage_collect</i>	Suppress normal garbage collection during backup; improves performance on some databases; corresponds to gbak -garbage_collect	—	Bit

TABLE 12.3 Services API database backup arguments

Argument	Purpose	Argument length	Argument value
<code>isc_spb_bkp_old_descriptions</code>	Output metadata in pre-4.0 format; corresponds to gbak -old_descriptions	—	Bit
<code>isc_spb_bkp_non_transportable</code>	Output backup file format with non-XDR data format; improves space and performance by a negligible amount; corresponds to gbak -nt	—	Bit
<code>isc_spb_bkp_convert</code>	Convert external table data to internal tables; corresponds to gbak -convert	—	Bit

TABLE 12.3 Services API database backup arguments

EXAMPLE 12.4 Starting a database backup service in C/C++

```
char request[100],*x, *p = request;

/* Identify cluster */
*p++ = isc_action_svc_backup;

/* Argument for database filename */
*p++ = isc_spb_dbname;
ADD_SPB_LENGTH(p, strlen(argv[1]));
for (x = argv[1]; *x; ) *p++ = *x++;

/* Argument for backup output filename */
*p++ = isc_spb_bkp_file;
ADD_SPB_LENGTH(p, strlen(argv[2]));
for (x = argv[2]; *x; ) *p++ = *x++;

/* Argument to request verbose output */
*p++ = isc_spb_verbose;
```

```

if (isc_service_start(status,
    &service_handle,
    NULL,
    p - request,
    request))
{
    isc_print_status(status);
    isc_service_detach(status, service_handle);
    exit(1);
}

```

You can also restore database backup files to create a new `.gdb` file. The following table lists arguments to the cluster identifier `isc_action_svc_restore`:

Argument	Purpose	Argument length	Argument value
<code>isc_spb_bkp_file</code>	The path of the backup file name	2-byte length + string	String
<code>isc_spb_dbname</code>	Path of the primary file of the database, from the server's point of view; you can specify multiple database files	2-byte length + string	String
<code>isc_spb_res_length</code>	The length in pages of the restored database file; must not exceed 2 gigabytes; you must supply a length for each database file except the last	4 bytes	Unsigned long; pages in database file
<code>isc_spb_verbose</code>	If specified, the Services Manager prepares output to return via <code>isc_service_query()</code> ; corresponds to gbak -verbose	—	—
<code>isc_spb_res_buffers</code>	The number of default cache buffers to configure for attachments to the restored database; corresponds to gbak -buffers	4 bytes	Unsigned long; number of buffers
<code>isc_spb_res_page_size</code>	The page size for the restored database; corresponds to gbak -page_size		

TABLE 12.4 Services API database restore arguments

Argument	Purpose	Argument length	Argument value
<code>isc_spb_res_access_mode</code>	Set the access mode of the database; the next byte must be one of: <ul style="list-style-type: none"> • <code>isc_spb_prp_am_readonly</code> • <code>isc_spb_prp_am_readwrite</code> Corresponds to gbak -mode	1 byte	Byte
<code>isc_spb_options</code>	The following value is a bitmask of <code>isc_spb_res_xxxx</code> options below	4 bytes	Bitmask
<code>isc_spb_res_deactivate_idx</code>	Do not build user indexes during restore; corresponds to gbak -inactive	—	Bit
<code>isc_spb_res_no_shadow</code>	Do not recreate shadow files during restore; corresponds to gbak -kill	—	Bit
<code>isc_spb_res_no_validity</code>	Do not enforce validity conditions (for example, NOT NULL) during restore; corresponds to gbak -no_validity	—	Bit
<code>isc_spb_res_one_at_a_time</code>	Commit after completing restore of each table; corresponds to gbak -one_at_a_time	—	Bit
<code>isc_spb_res_replace</code>	Replace database, if one exists; corresponds to gbak -replace	—	Bit
<code>isc_spb_res_create</code>	Restore but do not overwrite an existing database; corresponds to gbak -create	—	Bit
<code>isc_spb_res_use_all_space</code>	Do not reserve 20% of each data page for future record versions; useful for read-only databases; corresponds to gbak -use_all_space	—	Bit

TABLE 12.4 Services API database restore arguments

EXAMPLE 12.5 Starting a database restore service in C/C++

```
char request[100], *x, *p = request;
unsigned long options;

/* Identify cluster */
*p++ = isc_action_svc_restore;
```

```

/* Arguments for backup filenames */
for (i = 1; argc > 1; --argc; ++i)
{
    *p++ = isc_spb_bkp_file;
    ADD_SPB_LENGTH(p, strlen(argv[i]));
    for (x = argv[i]; *x; ) *p++ = *x++;
}

/* Argument for database filename */
*p++ = isc_spb_db_name;
ADD_SPB_LENGTH(p, strlen(argv[i]));
for (x = argv[i]; *x; ) *p++ = *x++;

/* Argument to request verbose output */
*p++ = isc_spb_verbose;

/* Argument to specify restore options */
*p++ = isc_spb_options;
options = isc_spb_res_create;
ADD_SPB_NUMERIC(p, options);

if (isc_service_start(status,
    &service_handle,
    NULL,
    p - request,
    request))
{
    isc_print_status(status);
    isc_service_detach(status, service_handle);
    exit(1);
}

```

Setting database properties

You can configure the properties of local or remote databases using the cluster identifier `isc_action_svc_properties`. This functionality corresponds to several of the options of the `gfix` command-line utility.

The following table lists arguments to `isc_action_svc_properties`:

Argument	Purpose	Argument length	Argument value
<code>isc_spb_dbname</code>	Path of the primary file of the database, from the server's point of view	2-byte length + string	String
<code>isc_spb_prp_page_buffers</code>	Set the default number of cache buffers to the specified number; corresponds to <code>gfix -buffers</code>	4 bytes	Unsigned long
<code>isc_spb_prp_sweep_interval</code>	Set the sweep interval to the specified number; specify zero to disable sweeping; corresponds to <code>gfix -housekeeping</code>	4 bytes	Unsigned long
<code>isc_spb_prp_shutdown_db</code>	Shuts down the database when: <ul style="list-style-type: none"> • There are no connections to the database, or • At the end of the timeout period you specify Corresponds to <code>gfix -shut -force n</code>	4 bytes	Unsigned long
<code>isc_spb_prp_deny_new_transactions</code>	Shuts down the database if there are no active transactions at the end of the timeout period you specify; deny new transactions during this timeout period; fail if there are active transactions at the end of the timeout period; corresponds to <code>gfix -shut -tran n</code>	4 bytes	Unsigned long
<code>isc_spb_prp_deny_new_attachments</code>	Shuts down the database if there are no active transactions at the end of the timeout period you specify; deny new database attachments during this timeout period; fail if there are active database attachments at the end of the timeout period; corresponds to <code>gfix -shut -attach n</code>	4 bytes	Unsigned long

TABLE 12.5 Services API database properties arguments

Argument	Purpose	Argument length	Argument value
<i>isc_spb_prp_reserve_space</i>	Configure the database to fill data pages when inserting new records, or reserve 20% of each page for later record deltas; the next byte must be one of: <ul style="list-style-type: none"> • <i>isc_spb_prp_res_use_full</i> • <i>isc_spb_prp_res</i> Corresponds to gfix -use	1 byte	Byte
<i>isc_spb_prp_write_mode</i>	Set the write mode for the database; the next byte must be one of: <ul style="list-style-type: none"> • <i>isc_spb_prp_wm_async</i> • <i>isc_spb_prp_wm_sync</i> Corresponds to gfix -write	1 byte	Byte
<i>isc_spb_prp_access_mode</i>	Set the access mode of the database; the next byte must be one of: <ul style="list-style-type: none"> • <i>isc_spb_prp_am_readonly</i> • <i>isc_spb_prp_am_readwrite</i> Corresponds to gfix -mode	1 byte	Byte
<i>isc_spb_prp_set_sql_dialect</i>	Set the SQL dialect for the database; value must be either 1 or 3	4 bytes	Unsigned long
<i>isc_spb_options</i>	The following value is a bitmask of <i>isc_spb_prp_XXXX</i> options below	4 bytes	Bitmask
<i>isc_spb_prp_activate</i>	Activate shadow file for use as a database; corresponds to gfix -activate	—	Bit
<i>isc_spb_prp_db_online</i>	Bring a shutdown database back online; corresponds to gfix -online	—	Bit

TABLE 12.5 Services API database properties arguments

Invoking database maintenance

This section describes how to use `isc_service_start()` to perform database validation, sweep garbage collection, and resolve limbo transactions. These tasks correspond to several of the options of the **gfix** command-line utility.

► *Invoking a database validation*

You can request a database validation with the cluster identifier `isc_action_svc_repair`. Database validation scans internal data structures for specific types of corruption. In some cases, the validation operation can repair corruption.

IMPORTANT The validation operation cannot guarantee to repair all cases of corruption. Do not rely on database validation as a disaster recovery policy in lieu of making regular backups of your database.

The following table lists arguments to `isc_action_svc_repair` to validate a database:

Argument	Purpose	Argument length	Argument value
<code>isc_spb_dbname</code>	Path of the primary file of the database, from the server's point of view	2-byte length + string	String
<code>isc_spb_options</code>	The following value is a bitmask of <code>isc_spb_rpr_xxxx</code> options below	4 bytes	Bitmask
<code>isc_spb_rpr_check_db</code>	Request read-only validation of the database, without correcting any problems; corresponds to gfix -no_update	—	Bit
<code>isc_spb_rpr_ignore_checksum</code>	Ignore all checksum errors; corresponds to gfix -ignore	—	Bit
<code>isc_spb_rpr_kill_shadows</code>	Remove references to unavailable shadow files; corresponds to gfix -kill	—	Bit
<code>isc_spb_rpr_mend_db</code>	Mark corrupted records as unavailable, so subsequent operations skip them; corresponds to gfix -mend	—	Bit

TABLE 12.6 Services API database validation arguments

Argument	Purpose	Argument length	Argument value
<i>isc_spb_rpr_validate_db</i>	Locate and release pages that are allocated but unassigned to any data structures; corresponds to gfix -validate	—	Bit
<i>isc_spb_rpr_full</i>	Check record and page structures, releasing unassigned record fragments; use with <i>isc_spb_rpr_validate_db</i> ; corresponds to gfix -full	—	Bit

TABLE 12.6 Services API database validation arguments

► *Invoking a database sweep*

You can invoke a database sweep with the cluster identifier *isc_action_svc_repair*. Sweeping attempts to scan the database for outdated record versions and mark them as free space. The following table lists arguments to *isc_action_svc_repair* to sweep a database:

Argument	Purpose	Argument length	Argument value
<i>isc_spb_dbname</i>	Path of the primary file of the database, from the server's point of view	2-byte length + string	String
<i>isc_spb_options</i>	The following value is a bitmask of <i>isc_spb_rpr_xxxx</i> options below	4 bytes	Bitmask
<i>isc_spb_rpr_sweep_db</i>	Request database sweep to mark outdated records as free space; corresponds to gfix -sweep	—	Bit

TABLE 12.7 Services API database sweep arguments

► *Resolving limbo transactions*

You can list and correct transactions in a limbo state using the cluster identifier *isc_action_svc_repair*.

Note Limbo transactions are the result of interruptions in the two-phase commit process of InterBase. Most client interfaces, including BDE and ODBC, do not exercise the two-phase commit or distributed transaction capabilities of InterBase, therefore applications using such client interfaces never create limbo transactions.

The following table lists arguments to `isc_action_svc_repair` to resolve limbo transactions in a database:

Argument	Purpose	Argument length	Argument value
<code>isc_spb_dbname</code>	Path of the primary file of the database, from the server's point of view	2-byte length + string	String
<code>isc_spb_rpr_commit_trans</code>	Request that the Services Manager commit the transactions that follow	—	—
<code>isc_spb_rpr_rollback_trans</code>	Request that the Services Manager roll back the transactions that follow	—	—
<code>isc_spb_rpr_recover_two_phase</code>	Request that the Services Manager use automatic two-phase commit recovery on the specified transactions	—	—
<code>isc_spb_tra_id</code>	Precedes a transaction ID number	4 bytes	Unsigned long

TABLE 12.8 Services API limbo transaction arguments

Requesting database and server status reports

This section describes how to request database statistics and the server error log.

► *Requesting database statistics*

You can use the cluster identifier `isc_action_svc_db_stats` to request the Service Manager prepare database statistics. This corresponds to the functionality of the `gstat` command-line utility. You can subsequently receive this information using `isc_service_query()` (see “**Querying service tasks**” on page 237). The following table lists arguments to `isc_action_svc_db_stats`:

Argument	Purpose	Argument length	Argument value
<code>isc_spb_dbname</code>	Path of the primary file of the database, from the server's point of view	2-byte length + string	String
<code>isc_spb_options</code>	The following value is a bitmask of <code>isc_spb_sts_xxxx</code> options below	4 bytes	Bitmask
<code>isc_spb_sts_data_pages</code>	Request statistics for user data pages; corresponds to <code>gstat -data</code>	—	Bit
<code>isc_spb_sts_db_log</code>	Request only the information in the database log pages; corresponds to <code>gstat -log</code>	—	Bit
<code>isc_spb_sts_hdr_pages</code>	Request only the information in the database header page; corresponds to <code>gstat -header</code>	—	Bit
<code>isc_spb_sts_idx_pages</code>	Request statistics for user index pages; corresponds to <code>gstat -index</code>	—	Bit
<code>isc_spb_sts_sys_relations</code>	Request statistics for system tables and indexes in addition to user tables and indexes; corresponds to <code>gstat -system</code>	—	Bit

TABLE 12.9 Services API status report arguments

► *Requesting the server log*

You can use the cluster identifier `isc_action_svc_get_ib_log` to request the Services Manager to return the contents of the `interbase.log` file from the server. There are no arguments for this cluster.

You can retrieve the text that the server manager returns by using `isc_service_query()`. See “**Querying service tasks**” on page 237.

Configuring users

You can use the Services API to display, add, delete, and modify users. This corresponds to the functionality of the command-line tool `gsec`.

► *Listing valid users in the security database*

The following table lists arguments to `isc_action_svc_display_users`:

Argument	Purpose	Argument length	Argument value
<code>isc_spb_sec_username</code>	Specify a single user by name for which the Services Manager should return information	2 byte length + string	String

TABLE 12.10 Services API display users arguments

To request the Services Manager to return information for all users in `isc4.gdb`, omit the `isc_spb_sec_username` argument.

You can retrieve the information that the server manager returns by using `isc_service_query()` with the cluster identifier `isc_info_svc_get_users`. See **“Querying using Services API: server configuration information” on page 229**.

► *Adding a user to the security database*

You can create a new user in `isc4.gdb` using the cluster identifier `isc_action_svc_add_user`. The following table lists arguments to this cluster:

Argument	Purpose	Argument length	Argument value
<code>isc_spb_sec_username</code>	User name to create in <code>isc4.gdb</code> ; maximum 31 characters; <i>mandatory argument</i>	2 byte length + string	String
<code>isc_spb_sec_password</code>	Password for the user; maximum 31 characters, only first 8 characters are significant; <i>mandatory argument</i>	2 byte length + string	String
<code>isc_spb_sec_firstname</code>	Optional first name of the person using this user name	2 byte length + string	String
<code>isc_spb_sec_middlename</code>	Optional middle name of the person using this user name	2 byte length + string	String

TABLE 12.11 Services API add user arguments

Argument	Purpose	Argument length	Argument value
<i>isc_spb_sec_lastname</i>	Optional last name of the person using this user name	2 byte length + string	String
<i>isc_spb_sec_userid</i>	Optional user ID number, defined in <i>/etc/passwd</i> , to assign to the user in <i>isc4.gdb</i> ; reserved for future implementation	4 bytes	Unsigned long
<i>isc_spb_sec_groupid</i>	Optional group ID number, defined in <i>/etc/group</i> , to assign to the user in <i>isc4.gdb</i> ; reserved for future implementation	4 bytes	Unsigned long
<i>isc_spb_sec_groupname</i>	Optional group name, as defined in <i>/etc/group</i> , to assign to the user in <i>isc4.gdb</i> ; reserved for future implementation	2 byte length + string	String
<i>isc_spb_sql_role_name</i>	Optional SQL role to adopt when administering users (reserved for future use)	2 byte length + string	String

TABLE 12.11 Services API add user arguments

► *Removing a user from the security database*

You can create a new user in *isc4.gdb* using the cluster identifier *isc_action_svc_delete_user*. The following table lists arguments to this cluster:

Argument	Purpose	Argument length	Argument value
<i>isc_spb_sec_username</i>	Name of user to delete from <i>isc4.gdb</i> ; <i>mandatory argument</i>	2 byte length + string	String
<i>isc_spb_sql_role_name</i>	Optional SQL role to adopt when administering users (reserved for future use)	2 byte length + string	String

TABLE 12.12 Services API remove user arguments

If you remove a user entry from *isc4.gdb*, no one can log in to any database on that server using that name. You must create a new entry for that name using *isc_action_svc_add_user*.

► *Modifying a user in the security database*

You can create a new user in **isc4.gdb** using the cluster identifier `isc_action_svc_modify_user`.

The arguments you can use in this cluster are the same as those you can use with `isc_action_svc_add_user`. You cannot change a user name, only associated properties of that user entry. Only properties you specify change. To remove a property, specify zero for the length and data of the property.

► *Deprecated use of InterBase 5 user functions*

The API functions `isc_add_user()`, `isc_delete_user()`, and `isc_modify_user()` are made obsolete by the introduction of the InterBase Services API. The new Services API functions are preferred over the version 5 user configuration functions in order to provide a consistent services mechanism, interface, and set of messages.

It is recommended that you use the Services API functions instead of the obsolete user configuration functions. The `isc_xxxx_user()` functions are still present in InterBase version 6 for backward compatibility, but they are likely to be removed from the product in a future release.

Administering software activation certificates

You can use the Services API to install or remove software activation certificates. Use the cluster identifiers `isc_action_svc_add_license` and `isc_action_svc_remove_license`, respectively.

The following table lists arguments to `isc_action_svc_add_license` and `isc_action_svc_remove_license`:

Argument	Purpose	Argument length	Argument value
<code>isc_spb_lic_key</code>	The key string identifying a software activation certificate	2 byte length + string	String
<code>isc_spb_lic_id</code>	The ID string for a software activation certificate (<code>isc_action_svc_add_license</code> only)	2 byte length + string	String

TABLE 12.13 Services API software activation certificate arguments

► *Listing software activation certificates*

You can use `isc_service_query()` with the `isc_info_get_license` cluster identifier to

See “**Querying using Services API: software activation certificates**” on page 227 for an example of retrieving the certificate information with `isc_service_query()`.

▶ *Adding a software activation certificate*

To add a software activation certificate, you must specify both the certificate ID and certificate key in the respective arguments to `isc_action_svc_add_license`.

▶ *Removing a software activation certificate*

To remove a software activation certificate, you need specify only the certificate key in the respective arguments to `isc_action_svc_add_license`.

▶ *Enabling changes to certificates*

Changes to the number of simultaneous users take effect immediately.

In order for other changes to the certificates to take effect, you must stop and restart the InterBase service. There is no Services API method in the current implementation to restart the InterBase service.

You can stop and start a service on the local Windows NT host programmatically using the Win32 API. You must be Administrator or a member of the Power Users NT group to start or stop a service.

On Superserver implementations of InterBase on UNIX, you must use `ibmgr` to shut down and restart `ibserver`. Classic InterBase servers read the `iblicense.dat` file each time an instance of the server starts.

Querying the Services Manager

You can use the Services API function `isc_service_query()` to request information from the Services Manager about the InterBase server environment. This section describes how to request and interpret data from `isc_service_query()`.

Blocking and specifying timeout

You can request output of a service task in progress using `isc_service_query()`. Your call to `isc_service_query()` does not return until either the request has completed, or the result buffer is full. If there are no data to return because a service task is in progress, the query waits for the task to complete. `isc_service_query()` blocks indefinitely, until output is available. This eliminates the need for polling.

You can supply to *isc_service_query()* an SPB item specifying a finite duration after which the call to *isc_service_query()* must return, even if output from the task is not yet available. Populate the SPB with the SPB version information, followed by the *isc_info_svc_timeout* cluster identifier, and a four-byte value specifying the number of seconds for the timeout.

This is the only useful SPB cluster for *isc_service_query()* in the current implementation.

Services API query example

In this chapter, a complete C/C++ code sample shows use of the *isc_service_query()* function. The sample is split into several parts, to illustrate query items described in later sections. The code sample assumes that you have successfully attached to a Services Manager (see **“Attaching to the Services Manager with *isc_service_attach()*” on page 202**) and that you have a valid service handle.

The first part of the example shows how to set up the request buffer and invoke *isc_service_query()*.

EXAMPLE 12.6 Querying using Services API: setup and invoke query

```
char spb_buffer[6], *spb = spb_buffer;
char request_buffer[] = {
    isc_info_svc_server_version,
    isc_info_svc_implementation,
    isc_info_svc_get_licensed_users,
    isc_info_svc_user_dbpath,
    isc_info_svc_get_env,
    isc_info_svc_get_env_lock,
    isc_info_svc_get_env_msg,
    isc_info_svc_get_license,
    isc_info_svc_svr_db_info,
    isc_info_svc_version,
    isc_info_svc_get_config};
char result_buffer[1024], *p = result_buffer;

*spb++ = isc_info_svc_timeout;
ADD_SPB_NUMERIC(spb, 60); /* 1 minute timeout */
```

```

if (isc_service_query (
    status,
    &service_handle,
    NULL,
    spb - spb_buffer, spb_buffer,
    sizeof(request_buffer), request_buffer,
    sizeof(result_buffer), result_buffer))
{
    isc_print_status(status);
    isc_service_detach(status, &svc_handle);
    return;
}

do
{
    switch (*p++)
    {
        . . .

```

The code sample is continued in later examples.

Using result buffers

The Services API uses a buffer structured similarly to the SPB for *isc_service_query()* to specify tasks and options for the Services Manager. This is called the *request buffer*. You supply clusters of parameters and arguments in the request buffer. The Services Manager supplies the data you requested by specifying these arguments.

isc_service_query() uses another structured buffer to return requested data. This is called the *result buffer*. The Services Manager stores data in this buffer. You write code in your application to scan the buffer after *isc_service_query()* returns, and interpret the data based on the single-byte cluster identifiers at the start of each cluster.

The cluster identifiers are used both for requesting data in the request buffer, and for identifying clusters of returned data in the result buffer. When you add these identifiers to the request buffer, you specify only the identifier name in the request buffer, not the identifiers for any arguments. The Services Manager returns argument identifiers and data in the result buffer.

When you interpret the identifiers in the result buffer, clusters include associated data. The data that follow the cluster identifier are specific to the cluster type. Some clusters have a fixed length value following the identifier, for example numeric values are always returned as 4-byte long integers. Other clusters identifiers are followed by a 2-byte short integer, which specifies the length of the subsequent string. Still other cluster identifiers are followed by a series of argument identifiers with fixed or variable length data.

If the data that the Server Manager returns exceed the size of the result buffer you supply, *isc_service_query()* fills the buffer as much as possible, and includes *isc_info_truncated* as the last cluster identifier. This indicates that the result buffer was too small to contain all the resulting output of the service query. To receive the entire buffer, you must call *isc_service_query()* again with a larger buffer. The Services Manager starts over from the beginning of the output; you must provide a buffer that is large enough to hold the entire output.

EXAMPLE 12.7 Querying using Services API: handling a truncated result

```

. . .
case isc_info_truncated:
    printf ("Buffer Truncated\n");
    /* you should increase the buffer size and retry the query */
    break;
. . .

```

For output that is typically very lengthy, such as the output of a database backup task, the Services Manager needs to return a volume of text data. You can use the request item *isc_info_svc_line* to request successive lines of the text result, or you can use *isc_info_svc_to_eof* to request the entire text output in one query. See **“Querying service tasks” on page 237**.

Querying server configuration

You can use the following items with `isc_service_query()` to request information about the InterBase server configuration.

Server configuration items	Purpose	Return length	Return value
<code>isc_info_svc_version</code>	The version of the Services Manager	4 bytes	Unsigned long
<code>isc_info_svc_server_version</code>	The version of the InterBase server	2-byte length + string	String
<code>isc_info_svc_implementation</code>	The implementation string, or platform, of the server; for example, <i>InterBase/Sun4</i>	2-byte length + string	String
<code>isc_info_svc_get_license</code>	All software activation certificate IDs and keys currently enabled on the server	See below	See below
<code>isc_info_svc_get_license_mask</code>	A bitmask representing the software activation certificate options currently enabled on the server; reserved for future implementation	4 bytes	Bitmask
<code>isc_info_svc_capabilities</code>	A bitmask representing the capabilities currently enabled on the server; reserved for future implementation	4 bytes	Bitmask
<code>isc_info_svc_get_config</code>	The parameters and values in the ibconfig file on the server (isc_config on UNIX and Linux)	See below	See below
<code>isc_info_svc_get_env</code>	The location of the InterBase root directory on the server; this is the value of the <code>\$INTERBASE</code> system environment variable, or the contents of the registry key	2-byte length + string	String
<code>isc_info_svc_get_env_lock</code>	The location of the InterBase lock manager file on the server; this is the value of the <code>\$INTERBASE_LCK</code> system environment variable, or by default <code>\$INTERBASE/serverhostname.lck</code>	2-byte length + string	String
<code>isc_info_svc_get_env_msg</code>	The location of the InterBase message file on the server; this is the value of the <code>\$INTERBASE_MSG</code> system environment variable, or by default <code>\$INTERBASE/interbase.msg</code>	2-byte length + string	String

TABLE 12.14 Services API server configuration query items

EXAMPLE 12.8 Querying using Services API: Services Manager version

```

. . .
case isc_info_svc_version:
{
    unsigned long svcversion;
    p += sizeof (unsigned short);
    svcversion = (unsigned long)
        isc_vax_integer (p, sizeof(unsigned long));
    printf ("Service Manager Version: %d\n", svcversion);
    p += sizeof (unsigned long);
    break;
}
. . .

```

EXAMPLE 12.9 Querying using Services API: server version

```

. . .
case isc_info_svc_server_version:
{
    path_length = (unsigned short)
        isc_vax_integer (p, sizeof(unsigned short));
    p += sizeof (unsigned short);
    buffer = (char*) malloc (path_length);
    strncpy (buffer, p, path_length);
    buffer [path_length] = '\0';
    printf ("Server version: %s\n", buffer);
    p += path_length;
    break;
}
. . .

```

EXAMPLE 12.10 Querying using Services API: server implementation

```

. . .
case isc_info_svc_implementation:
{
    path_length = (unsigned short)
        isc_vax_integer (p, sizeof(unsigned short));
    p += sizeof (unsigned short);
    buffer = (char*) malloc (path_length);
    strncpy (buffer, p, path_length);

```

```

    buffer [path_length] = '\0';
    printf ("Server implementation: %s\n", buffer);
    p += path_length;
    break;
}
. . .

```

EXAMPLE 12.11 Querying using Services API: location of the server root directory

```

. . .
case isc_info_svc_get_env:
{
    path_length = (unsigned short)
        isc_vax_integer (p, sizeof(unsigned short));
    p += sizeof (unsigned short);
    buffer = (char*) malloc (path_length);
    strncpy (buffer, p, path_length);
    buffer [path_length] = '\0';
    printf ("Value of $INTERBASE: %s\n", buffer);
    free(buffer);
    p += path_length;
    break;
}
. . .

```

EXAMPLE 12.12 Querying using Services API: location of the server lock file

```

. . .
case isc_info_svc_get_env_lock:
{
    path_length = (unsigned short)
        isc_vax_integer (p, sizeof(unsigned short));
    p += sizeof (unsigned short);
    buffer = (char*) malloc (path_length);
    strncpy (buffer, p, path_length);
    buffer [path_length] = '\0';
    printf ("Path to <hostname>.lck: %s\n", buffer);
    free(buffer);
    p += path_length;
    break;
}
. . .

```

EXAMPLE 12.13 Querying the location of the message file using the Services API

```

. . .
case isc_info_svc_get_env_msg:
{
    path_length = (unsigned short)
        isc_vax_integer (p, sizeof(unsigned short));
    p += sizeof (unsigned short);
    buffer = (char*) malloc (path_length);
    strncpy (buffer, p, path_length);
    buffer [path_length] = '\0';
    printf ("Path to INTERBASE.MSG: %s\n", buffer);
    p += path_length;
    break;
}
. . .

```

► *Additional data for server configuration*

SOFTWARE ACTIVATION CERTIFICATES

The *isc_info_svc_get_license* result buffer item returns multiple sets of data as arguments. For each software activation certificate in the file **ib_license.dat** on the server, this cluster returns the ID and key strings. If there are multiple certificates installed on the server, the return buffer contains multiple pairs of ID and key strings. The contents of the buffer end when a cluster is identified with the *isc_info_flag_end* value. The following table describes the cluster identifiers for the certificate information.

Argument	Purpose	Return length	Return value
<i>isc_spb_lic_id</i>	The ID string for a software activation certificate	2-byte length + string	String
<i>isc_spb_lic_key</i>	The corresponding Key string for a software activation certificate	2-byte length + string	String
<i>isc_info_flag_end</i>	Signals the end of arguments to <i>isc_info_svc_get_license</i>	—	—

TABLE 12.15 Services API software activation certificate arguments

EXAMPLE 12.14 Querying using Services API: software activation certificates

```

. . .

```

```

case isc_info_svc_get_license:
{
    printf ("Software activation certificates:\n");
    do {
        switch (*p++)
        {
            case isc_spb_lic_key:
            {
                path_length = (unsigned short)
                    isc_vax_integer (p, sizeof(unsigned short));
                p += sizeof (unsigned short);
                buffer = (char*) malloc (path_length);
                strncpy (buffer, p, path_length);
                buffer [path_length] = '\0';
                printf ("\tLicense Key: %s\n", buffer);
                free(buffer);
                p += path_length;
                break;
            }
            case isc_spb_lic_id:
            {
                path_length = (unsigned short)
                    isc_vax_integer (p, sizeof(unsigned short));
                p += sizeof (unsigned short);
                buffer = (char*) malloc (path_length);
                strncpy (buffer, p, path_length);
                buffer [path_length] = '\0';
                printf ("\tLicense ID: %s\n", buffer);
                free(buffer);
                p += path_length;
                break;
            }
        }
    } while (*p != isc_info_flag_end);
    break;
}
. . .

```

SERVER CONFIGURATION PROPERTIES

You can request the Services Manager to report the contents of the InterBase configuration file on the server. This file is in the installation directory of InterBase, and is named **ibconfig** on Windows 95/98/NT and NetWare, and **isc_config** on UNIX and Linux.

The result buffer cluster consists of the *isc_info_svc_get_config* identifier, followed by a two-byte number of data. The data follow as pairs of single-byte configuration entry identifiers and four-byte values. Configuration entries with string values, such as *TMP_DIRECTORY*, are not currently supported by this cluster.

Some of the configuration items are relevant only on specific platforms. The Services Manager returns only configuration data that are relevant to the respective server platform that runs the Services Manager.

The Services Manager does not return values for configuration items that are set to their default value.

EXAMPLE 12.15 Querying using Services API: server configuration information

```

. . .
case isc_info_svc_get_config:
{
    unsigned short chTmp = 0, key;
    unsigned long len = 0, ulConfigInfo;

    printf ("Configuration Settings:\n");
    len = (unsigned short)
        isc_vax_integer(p, sizeof(unsigned short));
    p += sizeof(unsigned short);
    for (chTmp = 0; chTmp < len; chTmp++)
    {
        key = p[chTmp];
        ulConfigInfo = (unsigned long)
            isc_vax_integer(p+ chTmp + 2, p[chTmp+1]);
        switch (key)
        {
            case ISCCFG_LOCKMEM_KEY:
                printf ("\tLock mem: %d\n", ulConfigInfo);
                break;
            case ISCCFG_LOCKSEM_KEY:
                printf ("\tLock Semaphores: %d\n", ulConfigInfo);
                break;
            case ISCCFG_LOCKSIG_KEY:
                printf ("\tLock sig: %d\n", ulConfigInfo);

```

```

        break;
    case ISCCFG_EVTMEM_KEY:
        printf ("\tEvent mem: %d\n", ulConfigInfo);
        break;
    case ISCCFG_PRIORITY_KEY:
        printf ("\tPriority: %d\n", ulConfigInfo);
        break;
    case ISCCFG_MEMMIN_KEY:
        printf ("\tMin memory: %d\n", ulConfigInfo);
        break;
    case ISCCFG_MEMMAX_KEY:
        printf ("\tMax Memory: %d\n", ulConfigInfo);
        break;
    case ISCCFG_LOCKORDER_KEY:
        printf ("\tLock order: %d\n", ulConfigInfo);
        break;
    case ISCCFG_ANYLOCKMEM_KEY:
        printf ("\tAny lock mem: %d\n", ulConfigInfo);
        break;
    case ISCCFG_ANYLOCKSEM_KEY:
        printf ("\tAny lock semaphore: %d\n",
            ulConfigInfo);
        break;
    case ISCCFG_ANYLOCKSIG_KEY:
        printf ("\tany lock sig: %d\n", ulConfigInfo);
        break;
    case ISCCFG_ANYEVNTMEM_KEY:
        printf ("\tany event mem: %d\n", ulConfigInfo);
        break;
    case ISCCFG_LOCKHASH_KEY:
        printf ("\tLock hash: %d\n", ulConfigInfo);
        break;
    case ISCCFG_DEADLOCK_KEY:
        printf ("\tDeadlock: %d\n", ulConfigInfo);
        break;
    case ISCCFG_LOCKSPIN_KEY:
        printf ("\tLock spin: %d\n", ulConfigInfo);
        break;
    case ISCCFG_CONN_TIMEOUT_KEY:
        printf ("\tConn timeout: %d\n", ulConfigInfo);
        break;
    case ISCCFG_DUMMY_INTRVL_KEY:

```

```

        printf ("\tDummy interval: %d\n", ulConfigInfo);
        break;
    case ISCCFG_IPCMAP_KEY:
        printf ("\tMap size: %d\n", ulConfigInfo);
        break;
    case ISCCFG_DBCACHE_KEY:
        printf ("\tCache size: %d\n", ulConfigInfo);
        break;
    }
    chTmp += p[chTmp+1] + 1;
}
break;
}
. . .

```

Querying security configuration

You can use the following items with *isc_service_query()* to request information related to InterBase server security and user access.

Security configuration items	Purpose	Return length	Return value
<i>isc_info_svc_get_licensed_users</i>	The number of users permitted by the governor on the server	4 bytes	Unsigned long
<i>isc_info_svc_user_dbpath</i>	The path to the security database on the server; for example, <i>/usr/interbase/isc4.gdb</i>	2-byte length + string	String
<i>isc_info_svc_get_users</i>	User information from the security database	See below	See below
<i>isc_info_svc_svr_db_info</i>	The number of database attachments and databases currently active on the server	See below	See below

TABLE 12.16 Services API security configuration query items

EXAMPLE 12.16 Querying using Services API: number of licensed users

```

. . .
case isc_info_svc_get_licensed_users:
{
    unsigned long nUsers;
    p+= sizeof (unsigned short);
}

```

```

nUsers = (unsigned long)
    isc_vax_integer(p, sizeof (unsigned long));
printf ("Number of licensed users: %d\n", nUsers);
p += sizeof(unsigned long);
break;
}
. . .

```

EXAMPLE 12.17 Querying using Services API: location of the security database

```

. . .
case isc_info_svc_user_dbpath:
{
    path_length = (unsigned short)
        isc_vax_integer (p, sizeof(unsigned short));
    p += sizeof (unsigned short);
    buffer = (char*) malloc (path_length);
    strncpy (buffer, p, path_length);
    buffer [path_length] = '\0';
    printf ("Path to ISC4.GDB: %s\n", buffer);
    p += path_length;
    break;
}
. . .

```

► *Additional data for security information*

The `isc_info_svc_get_users` result item returns multiple sets of data. There might be multiple users to report, so the result buffer might contain multiple clusters. The contents of the buffer end when a cluster is identified with the `isc_info_flag_end` value. The following table describes the cluster identifiers for the user information.

Argument	Purpose	Return length	Return value
<code>isc_spb_username</code>	The user ID from <code>isc4.gdb</code>	2-byte length + string	String
<code>isc_spb_firstname</code>	The first name associated with the user ID	2-byte length + string	String
<code>isc_spb_middlename</code>	The middle name associated with the user ID	2-byte length + string	String

TABLE 12.17 Services API user information arguments

Argument	Purpose	Return length	Return value
<i>isc_spb_lastname</i>	The last name associated with the user ID	2-byte length + string	String
<i>isc_spb_userid</i>	The user ID number, defined in <i>/etc/passwd</i> , to assign to the user in <i>isc4.gdb</i> ; relevant only on UNIX or Linux servers	4 bytes	Unsigned long
<i>isc_spb_groupid</i>	The group ID number, defined in <i>/etc/group</i> , to assign to the user in <i>isc4.gdb</i> ; relevant only on UNIX or Linux servers	4 bytes	Unsigned long
<i>isc_info_flag_end</i>	Signals the end of arguments to <i>isc_info_svc_get_users</i>	—	—

TABLE 12.17 Services API user information arguments

EXAMPLE 12.18 Querying using Services API: users configured on the server

```

. . .
case isc_info_svc_get_users:
{
ISC_USHORT len, loop;
ISC_ULONG id;
char buffer[50], *buf = buffer;

loop = (ISC_USHORT)
isc_vax_integer (p, sizeof (ISC_USHORT));
p += sizeof (ISC_USHORT);

while (*p != isc_info_end)
{
switch (*p++)
{
case isc_spb_sec_username:
len = (ISC_USHORT)
isc_vax_integer(p, sizeof(ISC_USHORT));
p += sizeof (ISC_USHORT);
strncpy (buf, p, len);
p += len;
buffer[len] = 0;
printf ("Username: %s\n", buffer);
loop -= (len + sizeof(ISC_USHORT)+1);
}
}
}

```

```

        break;

    case isc_spb_sec_firstname:
        len = (ISC_USHORT)
            isc_vax_integer(p, sizeof(ISC_USHORT));
        p += sizeof (ISC_USHORT);
        strncpy (buf, p, len);
        p += len;
        buffer[len] = 0;
        printf ("Firstname: %s\n", buffer);
        loop -= (len + sizeof(ISC_USHORT)+1);
        break;

    case isc_spb_sec_middlename:
        len = (ISC_USHORT)
            isc_vax_integer(p, sizeof(ISC_USHORT));
        p += sizeof (ISC_USHORT);
        strncpy (buf, p, len);
        p += len;
        buffer[len] = 0;
        printf ("Middlename: %s\n", buffer);
        loop -= (len + sizeof(ISC_USHORT)+1);
        break;

    case isc_spb_sec_lastname:
        len = (ISC_USHORT)
            isc_vax_integer(p, sizeof(ISC_USHORT));
        p += sizeof (ISC_USHORT);
        strncpy (buf, p, len);
        p += len;
        buffer[len] = 0;
        printf ("Lastname: %s\n", buffer);
        loop -= (len + sizeof(ISC_USHORT)+1);
        break;

    case isc_spb_sec_groupid:
        id = isc_vax_integer (p, sizeof (ISC_ULONG));
        p += sizeof (ISC_ULONG);
        printf ("Group ID: %d\n", id);
        loop -= (len + sizeof(ISC_ULONG)+1);
        break;

```

```

        case isc_spb_sec_userid:
            id = isc_vax_integer (p, sizeof (ISC_ULONG));
            p += sizeof (ISC_ULONG);
            printf ("User ID: %d\n", id);
            loop -= (len + sizeof(ISC_ULONG)+1);
            break;

        default:
            *x = *p;
            break;
    } /* end switch */
} /* end while */
break;
}
. . .

```

The *isc_info_svc_svr_db_info* result item returns multiple sets of data. There might be multiple active databases to report, so the result buffer might contain multiple clusters. The contents of the buffer end when a cluster is identified with the *isc_info_flag_end* value. The following table describes the cluster identifiers for the database connection information.

Argument	Purpose	Return length	Return value
<i>isc_spb_num_att</i>	The number of attachments currently in use on the server	4 bytes	Unsigned long
<i>isc_spb_num_db</i>	The number of databases currently in use on the server	4 bytes	Unsigned long
<i>isc_spb_dbname</i>	The name of one of the databases currently in use on the server; this item occurs once for each database in use	2-byte length + string	String
<i>isc_info_flag_end</i>	Signals the end of arguments to <i>isc_info_svc_svr_db_info</i>	—	—

TABLE 12.18 Services API database connection information arguments

EXAMPLE 12.19 Querying using Services API: database attachments

```

. . .
case isc_info_svc_svr_db_info:
{

```

```

printf ("Database information:\n");
do {
    switch (*p++)
    {
        case isc_spb_dbname:
        {
            /* Database names in use */
            path_length = (unsigned short)
                isc_vax_integer(p, sizeof(unsigned short));
            p += sizeof (unsigned short);
            buffer = (char*) malloc (path_length);
            strncpy (buffer, p, path_length);
            buffer [path_length] = '\0';
            printf ("Database in use: %s\n", buffer);
            p += path_length;
            break;
        }
        case isc_spb_num_att:
        {
            /* Num Attachments */
            unsigned long nAttachments;
            p+= sizeof (unsigned short);
            nAttachments = (unsigned long)
                isc_vax_integer(p, sizeof (unsigned long));
            printf ("\tNumber of attachments: %d\n",
                nAttachments);
            p += sizeof(unsigned long);
            break;
        }
        case isc_spb_num_db:
        {
            /* Num databases */
            unsigned long nDatabases;
            p+= sizeof (unsigned short);
            nDatabases = (unsigned long)
                isc_vax_integer(p, sizeof(unsigned long));
            printf ("\tNumber of databases: %d\n",
                nDatabases);
            p += sizeof(unsigned long);
            break;
        }
    }
}

```

```

        } while (*p != isc_info_flag_end);
        break;
    }
    . . .

```

EXAMPLE 12.20 Querying using Services API: end of example

```

        . . .
    }
} while (*p);

isc_service_detach(status, &service_handle);
}

```

Querying service tasks

Some service tasks return textual output. You can use the following items with *isc_service_query()* to request output of a service task. The tasks that generate output are those corresponding to the following Services API task items: *isc_action_svc_backup*, *isc_action_svc_restore*, *isc_action_svc_repair*, *isc_action_svc_db_stats*, *isc_action_svc_get_ib_log*, and *isc_action_get_users*.

Task result items	Purpose	Return length	Return value
<i>isc_info_svc_line</i>	One line of output from a service task	2-byte length + string	Line of text
<i>isc_info_svc_to_eof</i>	Multiple lines of output from a service task, up to the size of the result buffer	2-byte length + string	Buffer of text
<i>isc_info_svc_running</i>	Returns TRUE if a service task is already running on the server; used for a non-blocking check for a task in progress	4 bytes	Unsigned long; 1 or 0
<i>isc_info_svc_get_users</i>	See: “Additional data for security information” on page 232	q.v.	q.v.
<i>isc_info_svc_limbo_trans</i>	Limbo transaction information for unresolved distributed transactions	See below	See below

TABLE 12.19 Services API task query items

► *Additional data for task results*

The *isc_info_svc_limbo_trans* result item returns multiple sets of data. If there are multiple limbo transactions to report, the result buffer contains multiple clusters. The contents of the buffer end when a cluster is identified with the *isc_info_flag_end* value. The following table describes the cluster identifiers for the limbo transaction information.

Argument	Purpose	Return length	Return value
<i>isc_dpb_single_tra_id</i>	Specifies a single-database limbo transaction ID	4 bytes	Unsigned long
<i>isc_spb_multi_tra_id</i>	Specifies a multi-database limbo transaction ID	4 bytes	Unsigned long
<i>isc_spb_tra_host_site</i>	Hostname of the client host that initiated the transaction; for multi-database transactions only	2-byte length + string	String
<i>isc_spb_tra_advise</i>	Server recommendation for resolution of a limbo transaction; value is one of the following: <ul style="list-style-type: none"> • <i>isc_spb_tra_advise_commit</i> • <i>isc_spb_tra_advise_rollback</i> • <i>isc_spb_tra_advise_unknown</i> 	1 byte	Byte
<i>isc_spb_tra_remote_site</i>	Hostname of a server on which the transaction is in a limbo state	2-byte length + string	String
<i>isc_spb_tra_db_path</i>	Path of the primary file of the database in which the transaction is in a limbo state	2-byte length + string	String
<i>isc_spb_tra_state</i>	Introduces a byte specifying the state of the limbo transaction; value is one of the following: <ul style="list-style-type: none"> • <i>isc_spb_tra_state_limbo</i> • <i>isc_spb_tra_state_commit</i> • <i>isc_spb_tra_state_rollback</i> • <i>isc_spb_tra_state_unknown</i> (this state should never occur) 	1 byte	Byte
<i>isc_info_flag_end</i>	Signals the end of arguments to <i>isc_info_svc_limbo_trans</i>	—	—

TABLE 12.20 Services API limbo transaction arguments

Using the Services API with Delphi and C++Builder

InterBase Express™ components product provides a package of Data Access Components for the visual development tools Delphi and C++Builder. This package includes a component interface to the Services API described in this chapter. Refer to the *Developer's Guide* for documentation of the InterBase Express™ services components.

PART II

API Reference Guide



CHAPTER 13

API Function Reference

This chapter is an alphabetical reference for the InterBase API function calls. It provides tables that categorize calls by the tasks they perform, and then provides an alphabetical and detailed description of each call, including its syntax, arguments, examples of use, and cross references to related calls.

Function categories

There are eleven classes of InterBase API function calls:

- Array functions for handling arrays of data
- Blob functions for handling the InterBase Blob datatype
- Database functions for handling database requests
- Conversion functions for translating dates between InterBase format and UNIX format, and for reversing the byte-order of integers
- DSQL functions for handling SQL statements entered by users at run time
- Error-handling functions
- Event functions for registering interest in events posted by triggers and stored procedures in applications and for processing the event queue

- Information functions for retrieving information about databases, transactions, Blob data, and events
- Security functions for adding, deleting, and modifying user records in the password database
- Services functions for administering server and database properties
- Transaction functions for handling transactions in an application

Some functions, such as information calls, occur in more than one class.

The embedded installation functions are not counted as part of the InterBase client API. They typically aren't used in database client applications, but for product installation applications. See the *Developer's Guide* for information about creating custom installation applications using the embedded installation functions.

Array functions

The following table summarizes the InterBase API functions available for handling array data in an application:

Function name	Purpose
<i>isc_array_get_slice()</i>	Retrieve a specified part of an array field
<i>isc_array_lookup_bounds()</i>	Determine the dimensions of an array field
<i>isc_array_lookup_desc()</i>	Retrieve an array description
<i>isc_array_put_slice()</i>	Write a specified part of an array field
<i>isc_array_set_desc()</i>	Set an array description

TABLE 13.1 Array functions

Blob functions

The following table summarizes the InterBase API functions available for handling Blob data in an application:

Function name	Purpose
<i>isc_blob_default_desc()</i>	Set a default Blob description for dynamic access
<i>isc_blob_gen_bpb()</i>	Generate a Blob parameter buffer (BPB) for dynamic access
<i>isc_blob_info()</i>	Request information about a Blob field
<i>isc_blob_lookup_desc()</i>	Retrieve a Blob description
<i>isc_blob_set_desc()</i>	Set a Blob description
<i>isc_cancel_blob()</i>	Discard a Blob
<i>isc_close_blob()</i>	Close a Blob
<i>isc_create_blob2()</i>	Create a new Blob
<i>isc_get_segment()</i>	Retrieve a segment of Blob data
<i>isc_open_blob2()</i>	Open a Blob for read access
<i>isc_put_segment()</i>	Write a segment of Blob data

TABLE 13.2 Blob functions

Database functions

The following table summarizes the InterBase API functions available for handling database requests in an application:

Function name	Purpose
<i>isc_attach_database()</i>	Connect to an existing database
<i>isc_database_info()</i>	Request information about an attached database
<i>isc_detach_database()</i>	Disconnect from a database
<i>isc_drop_database()</i>	Delete an attached database and its associated files
<i>isc_expand_dpb()</i>	Build a database parameter buffer (DPB) dynamically
<i>isc_version()</i>	Retrieve database implementation number and on-disk structure (ODS) major and minor version numbers

TABLE 13.3 Database functions

Conversion functions

The following table summarizes the InterBase API functions available for translating between InterBase DATE, TIME, and TIMESTAMP format and the UNIX date format, and for reversing the byte-order of an integer:

Function name	Purpose
<i>isc_decode_sql_date()</i>	Translate a date from InterBase format to C <i>struct tm</i> format
<i>isc_encode_sql_date()</i>	Translate a date from C <i>struct tm</i> format to InterBase format
<i>isc_decode_sql_time()</i>	Translate a time from InterBase format to C <i>struct tm</i> format
<i>isc_encode_sql_time()</i>	Translate a time from C <i>tm</i> format to InterBase format
<i>isc_decode_timestamp()</i>	Translate a date and time from InterBase format to C <i>struct tm</i> format
<i>isc_encode_timestamp()</i>	Translate a date and time from C <i>struct tm</i> format to InterBase format
<i>isc_vax_integer()</i>	Reverse the byte-order of an integer

TABLE 13.4 Date and conversion functions

Note To provide backward compatibility, the *isc_encode_date()* and *isc_decode_date()* functions are still available. They behave exactly like *isc_encode_timestamp()* and *isc_decode_timestamp()*.

DSQL functions

The following table summarizes the InterBase API functions available for handling DSQL statements built or entered by users at run time:

Function name	Purpose
<i>isc_dsql_allocate_statement()</i>	Allocate a statement handle
<i>isc_dsql_alloc_statement2()</i>	Allocate a statement handle that is automatically freed on database detachment
<i>isc_dsql_describe()</i>	Fill in an <i>XSQLDA</i> with information about values returned by a statement
<i>isc_dsql_describe_bind()</i>	Fill in an <i>XSQLDA</i> with information about a statement's input parameters
<i>isc_dsql_execute()</i>	Execute a prepared statement
<i>isc_dsql_execute2()</i>	Execute a prepared statement returning a single set of values
<i>isc_dsql_execute_immediate()</i>	Prepare and execute a statement without return values for one-time use
<i>isc_dsql_exec_immed2()</i>	Prepare and execute a statement with a single set of return values for one-time use
<i>isc_dsql_fetch()</i>	Retrieve data returned by a previously prepared and executed statement
<i>isc_dsql_free_statement()</i>	Free a statement handle, or close a cursor associated with a statement handle
<i>isc_dsql_prepare()</i>	Prepare a statement for execution
<i>isc_dsql_set_cursor_name()</i>	Define a cursor name and associate it with a statement handle
<i>isc_dsql_sql_info()</i>	Request information about a prepared statement

TABLE 13.5 DSQL functions

Error-handling functions

The following table summarizes the InterBase API functions available for handling database error conditions an application:

Function name	Purpose
<i>isc_interprete()</i>	Capture InterBase error messages to a buffer
<i>isc_print_sqlerror()</i>	Display an SQL error message
<i>isc_print_status()</i>	Display InterBase error messages
<i>isc_sqlcode()</i>	Set the value of SQLCODE
<i>isc_sql_interprete()</i>	Capture an SQL error message to a buffer

TABLE 13.6 Error-handling functions

Event functions

The following table summarizes the InterBase API functions available for handling events in an application:

Function name	Purpose
<i>isc_cancel_events()</i>	Cancel interest in an event
<i>isc_event_block()</i>	Allocate event parameter buffers
<i>isc_event_counts()</i>	Get the change in values of event counters in the event array
<i>isc_que_events()</i>	Wait asynchronously until an event is posted
<i>isc_wait_for_event()</i>	Wait synchronously until an event is posted

TABLE 13.7 Event functions

Information functions

The following table summarizes the InterBase API functions available for reporting information about databases, transactions, and Blob data to a client application that requests it:

Function name	Purpose
<i>isc_blob_info()</i>	Request information about a Blob field
<i>isc_database_info()</i>	Request information about an attached database
<i>isc_dsql_sql_info()</i>	Request information about a prepared DSQL statement
<i>isc_transaction_info()</i>	Request information about a specified transaction
<i>isc_version()</i>	Retrieve database implementation number and on-disk structure (ODS) major and minor version numbers

TABLE 13.8 Information functions

Security functions

The following table summarizes the InterBase API functions available for adding, deleting, and modifying user records in the password database:

Function name	Purpose
<i>isc_add_user()</i>	Adds a user to the password database
<i>isc_delete_user()</i>	Deletes a user from the password database
<i>isc_modify_user()</i>	Modifies user information in the password database

TABLE 13.9 Security functions

Services functions

The following table summarizes the InterBase API functions available for programmatic control of server and database administration tasks:

Function name	Purpose
<i>isc_service_attach()</i>	Attach to the InterBase Services Manager facility; required before using any of the InterBase services
<i>isc_service_detach()</i>	End the attachment to the InterBase Services Manager
<i>isc_service_query()</i>	Request and retrieve information about the InterBase server to which the client is attached
<i>isc_service_start()</i>	Perform a service task on the InterBase server to which the client is attached

TABLE 13.10 Service functions

Transaction control functions

The following table summarizes the InterBase API functions available for controlling transactions in an application:

Function name	Purpose
<i>isc_commit_retaining()</i>	Commit a transaction, and start a new one using the original transaction's context
<i>isc_commit_transaction()</i>	Save a transaction's database changes, and end the transaction
<i>isc_prepare_transaction()</i>	Execute the first phase of a two-phase commit
<i>isc_prepare_transaction2()</i>	Execute the second phase of a two-phase commit
<i>isc_rollback_transaction()</i>	Undo a transaction's database changes, and end the transaction
<i>isc_start_multiple()</i>	Begin new transactions (used on systems that do not support a variable number of input arguments)
<i>isc_start_transaction()</i>	Begin new transactions
<i>isc_transaction_info()</i>	Request information about a specified transaction

TABLE 13.11 Transaction control functions

Using function definitions

Each function definition in this chapter includes the elements in the following table:

Element	Description
Title	Function name
Definition	Main purpose of function
Syntax	Diagram of the function and parameters
Parameters	Table describing each parameter
Description	Detailed information about using the function
Example	Example of using the function in a program
Return value	Description of possible values returned in the status vector, if any
See also	Cross references to other related functions

TABLE 13.12 Function description format

isc_add_user()

Adds a user record to the password database, **isc4.gdb**.

Note Use of this function is deprecated. It is replaced by a full featured Services API. See **Chapter 12: “Working with Services” on page 199** and the reference entry for **“isc_service_start()” on page 380**.

Syntax

```
ISC_STATUS isc_add_user(
    ISC_STATUS *status
    USER_SEC_DATA *user_sec_data);
```

Parameter	Type	Description
<i>status vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>user_sec_data</i>	<i>USER_SEC_DATA *</i>	Pointer to a struct that is defined in ibase.h

Description The three security functions, *isc_add_user()*, *isc_delete_user()*, and *isc_modify_user()* mirror functionality that is available in the **gsec** command-line utility. *isc_add_user()* adds a record to **isc4.gdb**, InterBase’s password database.

At a minimum, you must provide the user name and password. If the server is not local, you must also provide a server name and protocol. Valid choices for the protocol field are *sec_protocol_tcpip*, *sec_protocol_netbeui*, *sec_protocol_spx*, and *sec_protocol_local*.

InterBase reads the settings for the `ISC_USER` and `ISC_PASSWORD` environment variables if you do not provide a DBA user name and password.

The definition for the `USER_SEC_DATA` structure in `ibase.h` is as follows:

```
typedef struct {
    short  sec_flags;      /* which fields are specified */
    int    uid;           /* the user's id */
    int    gid;           /* the user's group id */
    int    protocol;     /* protocol to use for connection */
    char   *server;      /* server to administer */
    char   *user_name;   /* the user's name */
    char   *password;    /* the user's password */
    char   *group_name;  /* the group name */
    char   *first_name;  /* the user's first name */
    char   *middle_name; /* the user's middle name */
    char   *last_name;   /* the user's last name */
    char   *dba_user_name; /* the dba user name */
    char   *dba_password; /* the dba password */
} USER_SEC_DATA;
```

When you pass this structure to one of the three security functions, you can tell it which fields you have specified by doing a bitwise OR of the following values, which are defined in `ibase.h`:

<code>sec_uid_spec</code>	<code>0x01</code>
<code>sec_gid_spec</code>	<code>0x02</code>
<code>sec_server_spec</code>	<code>0x04</code>
<code>sec_password_spec</code>	<code>0x08</code>
<code>sec_group_name_spec</code>	<code>0x10</code>
<code>sec_first_name_spec</code>	<code>0x20</code>
<code>sec_middle_name_spec</code>	<code>0x40</code>
<code>sec_last_name_spec</code>	<code>0x80</code>
<code>sec_dba_user_name_spec</code>	<code>0x100</code>
<code>sec_dba_password_spec</code>	<code>0x200</code>

No bit values are available for user name and password, since they are required.

The following error messages exist for this function:

Code	Value	Description
<i>isc_username_too_long</i>	335544747	The user name passed in is greater than 31 bytes
<i>isc_password_too_long</i>	335544748	The password passed in is longer than 8 bytes
<i>isc_username_required</i>	335544749	The operation requires a user name
<i>isc_password_required</i>	335544750	The operation requires a password
<i>isc_bad_protocol</i>	335544751	The protocol specified is invalid
<i>isc_dup_username_found</i>	335544752	The user name being added already exists in the security database
<i>isc_username_not_found</i>	335544753	The user name was not found in the security database
<i>isc_error_adding_sec_record</i>	335544754	An unknown error occurred while adding a user
<i>isc_error_deleting_sec_record</i>	335544755	An unknown error occurred while deleting a user
<i>isc_error_modifying_sec_record</i>	335544756	An unknown error occurred while modifying a user
<i>isc_error_updating_sec_db</i>	335544757	An unknown error occurred while updating the security database

TABLE 13.13 Error messages for user security functions

Example The following example adds a user (“Socks”) to the password database, using the bitwise *OR* technique for passing values from the *USER_SEC_DATA* structure.

```
{
    ISC_STATUS status[20];
    USER_SEC_DATA sec;

    sec.server           = "kennel";
    sec.dba_user_name   = "sysdba";
    sec.dba_password    = "masterkey";
    sec.protocol        = sec_protocol_tcpip;
    sec.first_name      = "Socks";
    sec.last_name       = "Clinton";
    sec.user_name       = "socks";
    sec.password        = "2meow!"; /* Note: do not hardcode passwords
*/
    sec.sec_flags       = sec_server_spec
```

```

| sec_password_spec
| sec_dba_user_name_spec
| sec_dba_password_spec
| sec_first_name_spec
| sec_last_name_spec;
isc_add_user(status, &sec);
/* check status for errors */
if (status[0] == 1 && status[1])
{
    switch (status[1]) {
        case isc_username_too_long:
            printf("Security database cannot accept long user names\n");
            break;
        ...
    }
}
}

```

Return Value `isc_add_user()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. See the “Description” section for this function for a list of error codes. For more information about examining the status vector, see **Chapter 10**, “**Handling Error Conditions.**”

See Also **`isc_delete_user()`**, **`isc_modify_user()`**

isc_array_get_slice()

Retrieves data from an array column in a row returned by a SELECT.

Syntax `ISC_STATUS isc_array_get_slice(`
`ISC_STATUS *status_vector,`
`isc_db_handle *db_handle,`
`isc_tr_handle *trans_handle,`
`ISC_QUAD *array_id,`
`ISC_ARRAY_DESC *desc,`
`void *dest_array,`
`ISC_LONG *slice_length);`

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle *</i>	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> ; the handle identifies the database containing the array column <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>trans_handle</i>	<i>isc_tr_handle *</i>	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL
<i>array_id</i>	<i>ISC_QUAD *</i>	Internal identifier for the array; the array ID must be previously retrieved through API DSQL functions
<i>desc</i>	<i>ISC_ARRAY_DESC *</i>	Descriptor defining the array slice (entire array or subset) to be retrieved
<i>dest_array</i>	<i>void *</i>	Pointer to a buffer of length <i>slice_length</i> into which the array slice will be copied by this function
<i>slice_length</i>	<i>ISC_LONG *</i>	Length, in bytes, of the <i>dest_array</i> buffer

isc_array_get_slice() retrieves data from an array column of a table row using an array ID. You can either retrieve all the array elements in that column, or a subset of contiguous array elements, called a *slice*. The upper and lower boundaries in the *desc* structure specify which elements are to be retrieved.

InterBase copies the elements into the buffer, *dest_array*, whose size is specified by *slice_length*. This should be at least the expected length required for the elements retrieved. Before returning from *isc_array_get_slice()*, InterBase sets *slice_length* to the actual number of bytes copied.

Before calling *isc_array_get_slice()*, there are many operations you must do in order to fill in the array descriptor, *desc*, determine the appropriate internal array identifier, *array_id*, and fetch the rows whose array columns you want to access. For complete step-by-step instructions for setting up an array descriptor and retrieving array information, see **Chapter 8, “Working with Array Data.”**

Note *Never* execute a DSQL statement that tries to access array column data directly unless you are fetching only a single element. The way to access slices of array column data is to call `isc_array_get_slice()` or `isc_array_put_slice()`. The only supported array references in DSQL statements are ones that specify an entire array column (that is, just the column name) in order to get the internal identifier for the array, which is required by `isc_array_get_slice()` and `isc_array_put_slice()`, or single element references.

Example The following program operates on a table named PROJ_DEPT_BUDGET. This table contains the quarterly head counts allocated for each project in each department of an organization. Each row of the table applies to a particular department and project. The quarterly head counts are contained in an array column named QUARTERLY_HEAD_CNT. Each row has four elements in this column, one per quarter. Each element of the array is a number of type *long*.

The example below selects the rows containing 1994 information for the project named VBASE. For each such row, it retrieves and prints the department number and the data in the array column (that is, the quarterly head counts).

In addition to illustrating the usage of `isc_array_lookup_bounds()` and `isc_array_get_slice()`, the program shows data structure initializations and calls to the DSQL functions required to prepare and execute the SELECT statement, to obtain the *array_id* needed by `isc_array_get_slice()`, and to fetch the selected rows one by one.

```
#include <ibase.h>
#define Return_if_Error(stat) if (stat[0] == 1 && stat[1]) \
    { \
        isc_print_status(stat); \
        return(1); \
    }

char *sel_str =
    "SELECT dept_no, quarterly_head_cnt FROM proj_dept_budget \
     WHERE year = 1994 AND proj_id = 'VBASE'";

char dept_no[6];
long hcnt[4], tr_handle, database_handle, SQLCODE;
short len, i, flag0, flag1;
ISC_QUAD array_id;
ISC_ARRAY_DESC desc;
ISC_STATUS status_vector[20], fetch_stat;
isc_stmt_handle stmt = NULL;
XSQLDA *osqlda;
tr_handle = database_handle = 0L;
/* Attach to a database here--this code omitted for brevity */
/* Start a transaction here--this code omitted for brevity */
/* Set up the SELECT statement. */
```

isc_array_get_slice()

```
/* Allocate the output XSQLDA for holding the array data. */
osqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(2));
osqlda->sqln = 2;
osqlda->version = 1;
/* Allocate a statement handle. */
isc_dsqli_allocate_statement(
    status_vector,
    &database_handle,
    &stmt);
Return_if_Error(status_vector);
/* Prepare the query for execution. */
isc_dsqli_prepare(
    status_vector,
    &tr_handle,
    &stmt,
    0,
    sel_str,
    1,
    osqlda);
Return_if_Error(status_vector);
/* Set up an XSQLVAR structure to allocate space for each item
   to be retrieved. */
osqlda->sqlvar[0].sqldata = (char *) dept_no;
osqlda->sqlvar[0].sqltype = SQL_TEXT + 1;
osqlda->sqlvar[0].sqlind = &flag0;
osqlda->sqlvar[1].sqldata = (char *) &array_id;
osqlda->sqlvar[1].sqltype = SQL_ARRAY + 1;
osqlda->sqlvar[1].sqlind = &flag1;
/* Execute the SELECT statement. */
isc_dsqli_execute(
    status_vector,
    &tr_handle,
    &stmt,
    1,
    NULL);
Return_if_Error(status_vector);
/* Set up the array descriptor. */
isc_array_lookup_bounds(
    status_vector,
    &database_handle, /* Set by previous isc_attach_database() call. */
    &tr_handle, /* Set by previous isc_start_transaction() call. */
    "PROJ_DEPT_BUDGET", /* Table name. */
```

```

    "QUARTERLY_HEAD_CNT", /* Array column name. */
    &desc);
Return_if_Error(status_vector);

/* Fetch the head count for each department's four quarters. */
while ((fetch_stat = isc_dsqli_fetch(
    status_vector,
    &stmt,
    1,
    osqlda)) == 0)
{
    if (!flag1)
    {
        /* There is array data; get the current values. */
        len = sizeof(hcnt);
        /* Fetch the data from the array column into hcnt array. */
        isc_array_get_slice(
            status_vector,
            &database_handle,
            &tr_handle,
            &array_id,
            &desc,
            hcnt,
            &len);
        Return_if_Error(status_vector);

        /* Print department number and head counts. */
        dept_no[osqlda->sqlvar[0].sqlllen] = '\0';
        printf("Department #: %s\n\n", dept_no);
        printf("\tCurrent counts: %d %d %d %d\n",
            hcnt[0], hcnt[1], hcnt[2], hcnt[3]);
    }
};
}
if (fetch_stat != 100L)
{
    SQLCODE = isc_sqlcode(status_vector);
    isc_print_sqlerror(SQLCODE, status_vector);
    return(1);
}
}

```

isc_array_lookup_bounds()

Return Value `isc_array_get_slice()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to `isc_bad_stmt_handle`, `isc_bad_trans_handle`, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_array_lookup_bounds()`, `isc_array_lookup_desc()`, `isc_array_put_slice()`, `isc_array_set_desc()`, `isc_dsql_fetch()`, `isc_dsql_prepare()`**

isc_array_lookup_bounds()

Determines the datatype, length, scale, dimensions, and array boundaries for the specified array column in the specified table.

Syntax

```
ISC_STATUS isc_array_lookup_bounds(  
    ISC_STATUS *status_vector,  
    isc_db_handle *db_handle,  
    isc_tr_handle *trans_handle,  
    char *table_name,  
    char *column_name,  
    ISC_ARRAY_DESC *desc);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS</i> *	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle</i> *	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> ; the handle identifies the database containing the array column <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>trans_handle</i>	<i>isc_tr_handle</i> *	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL
<i>table_name</i>	<i>char</i> *	Name of the table containing the array column, <i>column_name</i> ; can be either null-terminated or blank-terminated
<i>column_name</i>	<i>char</i> *	Name of the array column; can be either null-terminated or blank-terminated
<i>desc</i>	<i>ISC_ARRAY_DESC</i> *	Pointer to a descriptor for the arrays that will be filled in by this function

Description *isc_array_lookup_bounds()* determines the datatype, length, scale, dimensions, and array boundaries for the elements in an array column, *column_name* in the table, *table_name*. It stores this information in the array descriptor, *desc*.

isc_array_lookup_bounds() sets a flag in the descriptor to zero. This specifies that the array should be accessed in future function calls in row-major order, the default. If an application requires column-major access, reset this flag to 1.

The array descriptor is used in subsequent calls to *isc_array_get_slice()* or *isc_array_put_slice()*.

For a detailed description of the array descriptor, see **Chapter 8, “Working with Array Data.”**

Note There are ways to fill in an array descriptor other than by calling *isc_array_lookup_bounds()*. You can also:

- Call *isc_array_lookup_desc()*. This is exactly the same as calling *isc_array_lookup_bounds()*, except that the former does not fill in information about the upper and lower bounds of each dimension.
- Call *isc_array_set_desc()* to initialize the descriptor from parameters you call it with, rather than accessing the database metadata.

- Set the descriptor fields directly. Note that *array_desc_dtype* must be expressed as one of the datatypes in the following table, and the parameters, *array_desc_field_name*, and *array_desc_relation_name*, must be null-terminated:

array_desc_dtype	Corresponding InterBase datatype
<i>blr_text</i>	CHAR
<i>blr_text2</i>	CHAR
<i>blr_short</i>	SMALLINT
<i>blr_long</i>	INTEGER
<i>blr_quad</i>	ISC_QUAD structure
<i>blr_float</i>	FLOAT
<i>blr_double</i>	DOUBLE PRECISION
<i>blr_sql_date</i>	DATE
<i>blr_sql_time</i>	TIME
<i>blr_timestamp</i>	TIMESTAMP
<i>blr_varying</i>	VARCHAR
<i>blr_varying2</i>	VARCHAR
<i>blr_blob_id</i>	ISC_QUAD structure
<i>blr_cstring</i>	NULL-terminated string
<i>blr_cstring2</i>	NULL-terminated string

TABLE 13.14 Datatypes for array descriptor fields

Example The following illustrates a sample call to *isc_array_lookup_bounds()*. More complete examples of accessing arrays are found in the example programs for *isc_array_get_slice()* and *isc_array_put_slice()*.

```
#include <ibase.h>

ISC_STATUS status_vector[20];
ISC_ARRAY_DESC desc;
char *str1 = "PROJ_DEPT_BUDGET";
char *str2 = "QUARTERLY_HEAD_CNT";
```

```

isc_array_lookup_bounds(
    status_vector,
    &database_handle, /* Set in previous isc_attach_database() call. */
    &tr_handle, /* Set in previous isc_start_transaction() call. */
    str1,
    str2,
    &desc);
if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}

```

Return Value `isc_array_lookup_bounds()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to `isc_bad_stmt_handle`, `isc_bad_trans_handle`, `isc_fld_not_def`, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_array_get_slice()`, `isc_array_lookup_desc()`, `isc_array_put_slice()`, `isc_array_set_desc()`**

isc_array_lookup_desc()

Determines the datatype, length, scale, and dimensions for all elements in the specified array column in the specified table.

Syntax

```
ISC_STATUS isc_array_lookup_desc(
    ISC_STATUS *status_vector,
    isc_db_handle *db_handle,
    isc_tr_handle *trans_handle,
    char *table_name,
    char *column_name,
    ISC_ARRAY_DESC *desc);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle *</i>	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> ; the handle identifies the database containing the array column <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>trans_handle</i>	<i>isc_tr_handle *</i>	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL
<i>table_name</i>	<i>char *</i>	Name of the table containing the array column <i>column_name</i> ; can be either null-terminated or blank-terminated
<i>column_name</i>	<i>char *</i>	Name of the array column; can be either null-terminated or blank-terminated
<i>desc</i>	<i>ISC_ARRAY_DESC *</i>	Pointer to an array descriptor that will be filled in by this function

Description *isc_array_lookup_desc()* determines the datatype, length, scale, and dimensions for the array column, *column_name*, in the table, *table_name*. It stores this information in the array descriptor, *desc*.

It also sets to 0 a flag in the descriptor. This specifies that the array is accessed in future function calls in row-major order, the default. If an application requires column-major access, reset this flag to 1.

The array descriptor is used in subsequent calls to `isc_array_get_slice()` or `isc_array_put_slice()`.

For a detailed description of the array descriptor, see **Chapter 8, “Working with Array Data.”**

Note There are ways to fill in an array descriptor other than by calling `isc_array_lookup_desc()`. You can also:

- Call `isc_array_lookup_bounds()`. This is like `isc_array_lookup_desc()`, except that `isc_array_lookup_bounds()` also fills in information about the upper and lower bounds of each dimension.
- Call `isc_array_set_desc()`, to initialize the descriptor from parameters you call it with, rather than accessing the database metadata.
- Set the descriptor fields directly. Note that `array_desc_dtype` must be expressed as one of the datatypes in the following table, and the parameters, `array_desc_field_name`, and `array_desc_relation_name`, must be null-terminated:

array_desc_dtype	Corresponding InterBase datatype
<code>blr_text</code>	CHAR
<code>blr_text2</code>	CHAR
<code>blr_short</code>	SMALLINT
<code>blr_long</code>	INTEGER
<code>blr_quad</code>	ISC_QUAD structure
<code>blr_float</code>	FLOAT
<code>blr_double</code>	DOUBLE PRECISION
<code>blr_sql_date</code>	DATE
<code>blr_sql_time</code>	TIME
<code>blr_timestamp</code>	TIMESTAMP

TABLE 13.15 Datatypes for array descriptor fields

array_desc_dtype	Corresponding InterBase datatype
<i>blr_varying</i>	VARCHAR
<i>blr_varying2</i>	VARCHAR
<i>blr_blob_id</i>	ISC_QUAD structure
<i>blr_cstring</i>	NULL-terminated string
<i>blr_cstring2</i>	NULL-terminated string

TABLE 13.15 Datatypes for array descriptor fields (continued)

Example The following illustrates a sample call to *isc_array_lookup_desc()*. More complete examples of accessing arrays are found in the example programs for *isc_array_get_slice()* and *isc_array_put_slice()*.

```
#include <ibase.h>
ISC_STATUS status_vector[20];
ISC_ARRAY_DESC desc;
char str1 = "PROJ_DEPT_BUDGET";
char str2 = "QUARTERLY_HEAD_CNT";

isc_array_lookup_desc(
    status_vector,
    &database_handle, /* Set in previous isc_attach_database() call. */
    &tr_handle, /* Set in previous isc_start_transaction() call. */
    str1,
    str2,
    &desc);
if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
};
```

Return Value *isc_array_lookup_desc()* returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to *isc_bad_stmt_handle*, *isc_bad_trans_handle*, *isc fld_not_def*, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_array_get_slice()`, `isc_array_lookup_bounds()`, `isc_array_put_slice()`, `isc_array_set_desc()`**

isc_array_put_slice()

Writes data into an array column.

Syntax

```
ISC_STATUS isc_array_put_slice(
    ISC_STATUS *status_vector,
    isc_db_handle *db_handle,
    isc_tr_handle *trans_handle,
    ISC_QUAD *array_id,
    ISC_ARRAY_DESC *desc,
    void *source_array,
    ISC_LONG *slice_length);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle *</i>	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> ; the handle identifies the database containing the array column <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>trans_handle</i>	<i>isc_tr_handle *</i>	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL
<i>array_id</i>	<i>ISC_QUAD *</i>	On input, NULL (if you are creating a new array), or the internal identifier for an array to be modified, as assigned by the InterBase engine. This internal identifier must have been determined by previous calls to DSQL functions. This function changes <i>array_id</i> to be the identifier for the array it creates or modifies (see below).

Parameter	Type	Description
<i>desc</i>	<i>ISC_ARRAY_DESC *</i>	Descriptor defining the array slice (entire array or subset) to be written to
<i>source_array</i>	<i>void *</i>	Pointer to a buffer of length <i>slice_length</i> , that contains the slice of data that will be copied to the array by this function
<i>slice_length</i>	<i>ISC_LONG *</i>	Length, in bytes, of the <i>source_array</i> buffer

Description *isc_array_put_slice()* writes data into an array column. You can either store into all the array elements in that column, or into an array *slice*, a subset of contiguous array elements. The boundaries passed to the function in the array descriptor, *desc*, specify which elements are to be stored into.

InterBase copies the elements from the buffer, *source_array*, whose size is specified by *slice_length*.

The array identifier (array ID), *array_id*, should be passed as NULL if you are calling *isc_array_put_slice()* to create a new array. If you are calling it to modify an existing array, then *array_id* should be the identifier of the array to be modified. This must have been determined by previous calls to DSQL functions.

When *isc_array_put_slice()* is called with an array ID of an existing array, it:

1. Creates a new array with the same dimensions, bounds, etc., as the specified array, and copies the existing array data to the new array.
2. Writes the data from the array buffer, *source_array*, to the new array (or slice of the array), per the bounds specified in the array descriptor, *desc*.
3. Returns in the same *array_id* variable the array ID of the new array.

When *isc_array_put_slice()* is called with a NULL array ID, it:

1. Creates a new empty array with dimensions, bounds, etc., as declared for the array column whose name and table name are specified in the array descriptor, *desc*.
2. Writes the data from the array buffer, *source_array*, to the new array (or slice of the array)
3. Returns in the *array_id* variable the array ID of the new array.

Note that in both cases, a new array is created, and its array ID is returned in the *array_id* variable. The array is temporary until an UPDATE or INSERT statement is executed to associate the array with a particular column of a particular row.

You can make a single call to `isc_array_put_slice()` to write all the data you wish to the array. Or, you can call `isc_array_put_slice()` multiple times to store data into various slices of the array. In this case, each call to `isc_array_put_slice()` after the first call should pass the array ID of the temporary array. When `isc_array_put_slice()` is called with the array ID of a temporary array, it copies the specified data to the specified slice of the temporary array (it will not create a new array), and it doesn't modify `array_id`.

Before calling `isc_array_put_slice()`, there are many operations you must do in order to fill in the array descriptor, `desc`, determine the appropriate internal array identifier, `array_id`, and fetch the rows whose array columns you want to access.

For complete step-by-step instructions for setting up an array descriptor and writing array information, see **Chapter 8, "Working with Array Data."**

Note Never execute a DSQL statement that tries to directly store data into an array column. The *only* way to access array column data is by calling `isc_array_get_slice()` or `isc_array_put_slice()`. The only supported array references in DSQL statements are ones that specify an entire array column (that is, just the column name) in order to get the internal identifier for the array, which is required by `isc_array_get_slice()` and `isc_array_put_slice()`.

Example The following program operates on a table named PROJ_DEPT_BUDGET. This table contains the quarterly head counts allocated for each project in each department of an organization. Each row of the table applies to a particular department and project. The quarterly head counts are contained in an array column named QUARTERLY_HEAD_CNT. Each table row has four elements in this column, one per quarter. Each element is a number of type *long*.

This program selects the rows containing 1994 information for the project named VBASE. For each such row, it calls `isc_array_get_slice()` to retrieve a slice of the array, the quarterly head counts for the last two quarters. It then increments each, and calls `isc_array_put_slice()` to store the updated values.

In addition to illustrating the usage of `isc_array_lookup_desc()`, `isc_array_get_slice()`, and `isc_array_put_slice()`, the program shows data structure initializations and calls to the DSQL functions required to prepare and execute the SELECT and UPDATE statements, to obtain the `array_id` needed by `isc_array_get_slice()` and `isc_array_put_slice()`, to fetch the selected rows one by one, and to update the array ID.

```
#include <ibase.h>

#define Return_if_Error(stat) if (stat[0] == 1 && stat[1]) \
    { \
        isc_print_status(stat); \
```

isc_array_put_slice()

```
        return(1); \
    }

char *sel_str =
    "SELECT dept_no, quarterly_head_cnt FROM proj_dept_budget \
    WHERE year = 1994 AND proj_id = 'VBASE'";
char *upd_str =
    "UPDATE proj_dept_budget SET quarterly_head_count = ? \
    WHERE CURRENT OF S";

char dept_no[6];
long fetch_stat, SQLCODE, hcnt[2];
short len, i, flag0, flag1, flag2;
ISC_QUAD array_id;
ISC_ARRAY_DESC desc;
ISC_STATUS status_vector[20];
isc_stmt_handle stmt = NULL;
isc_stmt_handle ustmt = NULL;
char *cursor = "S";
XSQLDA *osqlda, *isqlda;

/* Set up the SELECT statement. */

/* Allocate the output XSQLDA for holding the array data. */
osqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(2));
osqlda->sqln = 2;
osqlda->version = SQLDA_VERSION1;

/* Allocate a statement handle for the SELECT statement. */
isc_dsqli_allocate_statement(
    status_vector, &database_handle, &stmt);
Return_if_Error(status_vector);

/* Prepare the query for execution. */
isc_dsqli_prepare(
    status_vector,
    &tr_handle,
    &stmt,
    0,
    sel_str,
    1,
    osqlda);
```

```

Return_if_Error(status_vector);

/* Set up an XSQLVAR structure to allocate space for each item
   to be retrieved. */

osqlda->sqlvar[0].sqldata = (char *) dept_no;
osqlda->sqlvar[0].sqltype = SQL_TEXT + 1;
osqlda->sqlvar[0].sqlind = &flag0;

osqlda->sqlvar[1].sqldata = (char *) &array_id;
osqlda->sqlvar[1].sqltype = SQL_ARRAY + 1;
osqlda->sqlvar[1].sqlind = &flag1;

/* Execute the SELECT statement. */
isc_dsql_execute(
    status_vector,
    &tr_handle,
    &stmt,
    1,
    NULL);
Return_if_Error(status_vector);

/* Declare a cursor. */
isc_dsql_set_cursor_name(
    status_vector, &stmt, cursor, 0);
Return_if_Error(status_vector);

/* Set up the UPDATE statement. */

/* Allocate a statement handle for the UPDATE statement. */
isc_dsql_allocate_statement(
    status_vector, &database_handle, &ustmt);
Return_if_Error(status_vector);

/* Allocate the input XSQLDA. */
isqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(2));
isqlda->sqln = 1;
isqlda->version = SQLDA_VERSION1;

/* Prepare the UPDATE statement for execution. */
isc_dsql_prepare(
    status_vector,

```

isc_array_put_slice()

```
        &tr_handle,
        &ustmt,
        0,
        upd_str,
        1,
        NULL);
Return_if_Error(status_vector);

/* Initialize the input XSQLDA. */
isc_dsql_describe_bind(
    status_vector, &ustmt, 1, isqlda);
Return_if_Error(status_vector);

/* Set up the input sqldata and sqlind fields. */
isqlda->sqlvar[0].sqldata = (char *) &array_id;
isqlda->sqlvar[0].sqlind = &flag2;

/* Set up the array descriptor. */
isc_array_lookup_desc(
    status_vector,
    &database_handle, /* Set by previous isc_attach_database() call. */
    &tr_handle, /* Set by previous isc_start_transaction() call. */
    "PROJ_DEPT_BUDGET", /* Table name. */
    "QUARTERLY_HEAD_CNT", /* Array column name. */
    &desc);
Return_if_Error(status_vector);

/* Set the descriptor bounds to those of the slice to be updated, that
is, to those of the last two elements. Assuming the array column was
defined to contain 4 elements, with a lower bound (subscript) of 1 and
an upper bound of 4, the last two elements are at subscripts 3 and 4. */
desc->array_desc_bounds[0].array_bound_lower = 3;
desc->array_desc_bounds[0].array_bound_upper = 4;

/* Fetch and process the rows of interest. */
while ((fetch_stat = isc_dsql_fetch(
    status_vector, &stmt, 1, osqlda)) == 0)
{
    if (!flag1)
    {
        /* There is array data; get values for last two quarters. */
        len = sizeof(hcnt);
```

```

/* Fetch the data from the array slice into hcnt array. */
isc_array_get_slice(
    status_vector,
    &database_handle,
    &tr_handle,
    &array_id,
    &desc,
    hcnt,
    &len);
Return_if_Error(status_vector);

/* Add 1 to each count. */
for (i = 0; i < 2; i++)
    hcnt[i] = hcnt[i] + 1;

/* Save new values. */
isc_array_put_slice(
    status_vector,
    &database_handle,
    &tr_handle,
    &array_id,
    &desc,
    hcnt,
    &len);
Return_if_Error(status_vector);

/* Update the array ID. */
isc_dsqli_execute(
    status_vector, &tr_handle, &ustmt, 1, isqlda);
Return_if_Error(status_vector);

};
};
if (fetch_stat != 100L)
{
    SQLCODE = isc_sqlcode(status_vector);
    isc_print_sqlerror(SQLCODE, status_vector);
    return(1);
}

```

isc_array_set_desc()

Return Value `isc_array_put_slice()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to `isc_bad_stmt_handle`, `isc_bad_trans_handle`, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_array_get_slice()`, `isc_array_lookup_bounds()`, `isc_array_lookup_desc()`, `isc_array_set_desc()`, `isc_dsql_allocate_statement()`, `isc_dsql_describe_bind()`, `isc_dsql_execute()`, `isc_dsql_fetch()`, `isc_dsql_prepare()`, `isc_dsql_set_cursor_name()`**

`isc_array_set_desc()`

Initializes an array descriptor.

Syntax

```
ISC_STATUS isc_array_get_slice(  
    ISC_STATUS *status_vector,  
    char *table_name,  
    char *column_name,  
    short *sql_dtype,  
    short *sql_length,  
    short *dimensions,  
    ISC_ARRAY_DESC *desc);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>table_name</i>	<i>char *</i>	Blank- or null-terminated name of the table containing the array column, <i>column_name</i>
<i>column_name</i>	<i>char *</i>	Name of the array column; this may be either null-terminated or blank-terminated
<i>sql_dtype</i>	<i>short *</i>	Pointer to SQL datatype of the array elements
<i>sql_length</i>	<i>short *</i>	Pointer to length of each array element
<i>dimensions</i>	<i>short *</i>	Pointer to number of array dimensions
<i>desc</i>	<i>ISC_ARRAY_DESC *</i>	Array descriptor to be filled in by this function

Description `isc_array_set_desc()` initializes the array descriptor, `desc`, from the function parameters, `table_name`, `column_name`, `sql_dtype`, `sql_length`, and `dimensions`.

`isc_array_set_desc()` also sets to 0 a flag in the descriptor. This specifies that the array is accessed in future function calls in row-major order, the default. If an application requires column-major access, reset this flag to 1.

`table_name` and `column_name` can be either null-terminated or blank-terminated. The names stored in the descriptor will be null-terminated.

`sql_dtype` must be given as an SQL macro constant.

The array descriptor is used in subsequent calls to `isc_array_get_slice()` or `isc_array_put_slice()`.

For a detailed description of the array descriptor, see **Chapter 8, “Working with Array Data.”**

Note There are ways to fill in an array descriptor other than by calling `isc_array_set_desc()`. You can also:

- Call `isc_array_lookup_bounds()`. This function is similar to `isc_array_lookup_desc()`, except that `isc_array_lookup_bounds()` also fills in information about the upper and lower bounds of each dimension.
- Call `isc_array_lookup_desc()`. This function is similar to `isc_array_lookup_bounds()`, except that `isc_array_lookup_desc()` does not fill in information about the upper and lower bounds of each dimension.
- Set the descriptor fields directly. Note that `array_desc_dtype` must be expressed as one of the datatypes in the following table, and the parameters, `array_desc_field_name`, and `array_desc_relation_name`, must be null-terminated:

array_desc_dtype	Corresponding InterBase datatype
<code>blr_text</code>	CHAR
<code>blr_text2</code>	CHAR
<code>blr_short</code>	SMALLINT
<code>blr_long</code>	INTEGER
<code>blr_quad</code>	ISC_QUAD structure
<code>blr_float</code>	FLOAT

TABLE 13.16 Datatypes for array descriptor fields

array_desc_dtype	Corresponding InterBase datatype
<i>blr_double</i>	DOUBLE PRECISION
<i>blr_sql_date</i>	DATE
<i>blr_sql_time</i>	TIME
<i>blr_timestamp</i>	TIMESTAMP
<i>blr_varying</i>	VARCHAR
<i>blr_varying2</i>	VARCHAR
<i>blr_blob_id</i>	ISC_QUAD structure
<i>blr_cstring</i>	NULL-terminated string
<i>blr_cstring2</i>	NULL-terminated string

TABLE 13.16 Datatypes for array descriptor fields (*continued*)

Example The following illustrates a sample call to `isc_array_set_desc()`. More complete examples of accessing arrays are found in the example programs for `isc_array_get_slice()` and `isc_array_put_slice()`.

```
#include <ibase.h>
ISC_STATUS status_vector[20];
ISC_ARRAY_DESC desc;
short dtype = SQL_TEXT;
short len = 8;
short dims = 1;

isc_array_set_desc(
    status_vector,
    "TABLE1",
    "CHAR_ARRAY",
    &dtype,
    &len,
    &dims,
    &desc);

if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}
```

```
}

```

Return Value `isc_array_set_desc()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_array_get_slice()`, `isc_array_lookup_bounds()`, `isc_array_lookup_desc()`, `isc_array_put_slice()`**

isc_attach_database()

Attaches to an existing database.

Syntax

```
ISC_STATUS isc_attach_database(
    ISC_STATUS *status_vector,
    short db_name_length,
    char *db_name,
    isc_db_handle *db_handle,
    short parm_buffer_length,
    char *parm_buffer);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>db_name_length</i>	<i>short</i>	Number of bytes in <i>db_name</i> string; if 0, the string is assumed to be null-terminated
<i>db_name</i>	<i>char *</i>	Database name
<i>db_handle</i>	<i>isc_db_handle *</i>	Pointer to a database handle set by this function; It is recommended that you set <i>db_handle</i> to NULL before passing it to <i>isc_attach_database()</i>
<i>parm_buffer_length</i>	<i>short</i>	Number of bytes in the database parameter buffer (DPB)
<i>parm_buffer</i>	<i>char *</i>	Address of the DPB

isc_attach_database()

Description The `isc_attach_database()` function connects to an existing database to enable subsequent program access. It also optionally specifies various operational characteristics, such as a user name and password combination for access to a database on a remote server, or the number of database cache buffers to use. These optional characteristics are passed in a database parameter buffer (DPB) supplied and populated by the calling program, either through direct program construction, and by calling `isc_expand_dpb()` to build the DPB.

A program passes the name of the database file to which to attach in `db_name`. For programs not written in C, the program must also pass the length, in bytes, of `db_name` in the `db_name_length` parameter. C programs should pass a 0 length in this parameter.

If successful, `isc_attach_database()` assigns a unique ID to `db_handle`. Subsequent API calls use this handle to identify the database against which they operate.

When finished accessing a database, disconnect from the database with `isc_detach_database()`.

Example The following program fragment attaches to a database named **employee.db**. In the parameter buffer, it specifies a user name and password. These come from the contents of char * variables named `user_name` and `user_password`, respectively.

```
char dpb_buffer[256], *dpb, *p;
ISC_STATUS status_vector[20];
isc_db_handle handle = NULL;
short dpb_length;

/* Construct the database parameter buffer. */
dpb = dpb_buffer;
*dpb++ = isc_dpb_version1;

*dpb++ = isc_dpb_user_name;
*dpb++ = strlen(user_name);
for (p = user_name; *p;)
    *dpb++ = *p++;

*dpb++ = isc_dpb_password;
*dpb++ = strlen(user_password);
for (p = user_password; *p;)
    *dpb++ = *p++;

/* An alternate choice for the above construction is to call:
isc_expand_dpb(). */

dpb_length = dpb - dpb_buffer;
```

```

isc_attach_database(
    status_vector,
    0,
    "employee.db",
    &handle,
    dpb_length,
    dpb_buffer);
if (status_vector[0] == 1 && status_vector[1])
{
    /* An error occurred. */
    isc_print_status (status_vector);
    return(1);
}

```

Return Value `isc_attach_database()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_detach_database()`, `isc_expand_dpb()`**

For more information about creating and populating a DPB, see **“Creating and populating a DPB” on page 44**. For more information about attaching to a database, see **“Connecting to databases” on page 42**.

isc_blob_default_desc()

Loads a data structure with default information about a Blob, including its subtype, character set, and segment size.

Syntax

```
void isc_blob_default_desc(
    ISC_BLOB_DESC *desc,
    unsigned char *table_name,
    unsigned char *column_name);
```

Parameter	Type	Description
<i>desc</i>	<i>ISC_BLOB_DESC *</i>	Pointer to a Blob descriptor
<i>table_name</i>	<i>unsigned char *</i>	Table name
<i>column_name</i>	<i>unsigned char *</i>	Blob column name

Description *isc_blob_default_desc()* loads a Blob descriptor, *desc*, with the specified *table_name* and *column_name*, and the following default values prior to calling *isc_blob_gen_bpb()* to generate a Blob parameter buffer (BPB) for the Blob column being accessed:

- Subtype is set to TEXT.
- Character set is set to the default character set for the process or database.
- Segment size is set to 80 bytes.

isc_blob_default_desc() and three related functions, *isc_blob_gen_bpb()*, *isc_blob_lookup_desc()*, and *isc_blob_set_desc()*, provide dynamic access to Blob information. In particular, these functions can define and access information about a Blob for filtering purposes, such as character set information for text Blob data, and subtype information for text and non-text Blob data.

The following table lists the fields in the *desc* structure:

Parameter	Type	Description
<i>blob_desc_subtype</i>	<i>short</i>	Subtype of the Blob filter
<i>blob_desc_charset</i>	<i>short</i>	Character set being used
<i>blob_desc_segment_size</i>	<i>short</i>	Blob segment size
<i>blob_desc_field_name</i> [32]	<i>char</i>	Array containing the name of the Blob column
<i>blob_desc_relation_name</i> [32]	<i>char</i>	Array containing the name of the table in which the Blob is stored

TABLE 13.17 Blob descriptor fields

Example The following fragment loads the Blob descriptor with default information:

```
typedef struct
{
    short          blob_desc_subtype;
    short          blob_desc_charset;
    short          blob_desc_segment_size;
    unsigned char  blob_desc_field_name[32];
    unsigned char  blob_desc_relation_name[32];
    ISC_BLOB_DESC;
    isc_blob_default_desc(&desc, &relation, &field);
}
```

Return Value None.

See Also **isc_blob_gen_bpb()**, **isc_blob_lookup_desc()**, **isc_blob_set_desc()**

For more information about Blob descriptors, see **Chapter 7, “Working with Blob Data.”**

isc_blob_gen_bpb()

Generates a Blob parameter buffer (BPB) to allow dynamic access to Blob subtype and character set information.

Syntax

```
ISC_STATUS isc_blob_gen_bpb(
    ISC_STATUS *status_vector,
    ISC_BLOB_DESC *to_desc,
    ISC_BLOB_DESC *from_desc,
    unsigned short bpb_buffer_length,
    unsigned char *bpb_buffer,
    unsigned short *bpb_length);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>to_desc</i>	<i>ISC_BLOB_DESC *</i>	Pointer to the target Blob descriptor
<i>from_desc</i>	<i>ISC_BLOB_DESC *</i>	Pointer to the source Blob descriptor
<i>bpb_buffer_length</i>	<i>unsigned short</i>	Length of the BPB <i>bpb_buffer</i>
<i>bpb_buffer</i>	<i>unsigned char *</i>	Pointer to the BPB
<i>bpb_length</i>	<i>unsigned short *</i>	Pointer to the length of the data stored into the BPB

Description *isc_blob_gen_bpb()* generates a Blob parameter buffer (BPB) from subtype and character set information stored in the source Blob descriptor *from_desc* and the target (destination) Blob descriptor *to_desc*.

A BPB is needed whenever a filter will be used when writing to or reading from a Blob column. Two Blob descriptors are needed for filtering: one (*from_desc*) to describe the filter source data, and the other (*to_desc*) to describe the destination. The descriptors must have been previously created either directly, or via a call to *isc_blob_default_desc()*, *isc_blob_lookup_desc()*, or *isc_blob_set_desc()*.

The BPB generated by *isc_blob_gen_bpb()* is subsequently needed in calls to *isc_open_blob2()* or *isc_create_blob2()* if filtering will be utilized. For more information about the BPB, see **Chapter 7, “Working with Blob Data.”**

Example The following fragment generates the Blob descriptor:

```
isc_blob_gen_bpb(status, &to_desc, &from_desc, bpb_length, &buffer,
    &buf_length);
```

Return Value `isc_blob_gen_bpb()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_blob_default_desc()`, `isc_blob_lookup_desc()`, `isc_blob_set_desc()`, `isc_create_blob2()`, `isc_open_blob2()`**

isc_blob_info()

Returns information about an open Blob.

Syntax

```
ISC_STATUS isc_blob_info(
    ISC_STATUS *status_vector,
    isc_blob_handle *blob_handle,
    short item_list_buffer_length,
    char *item_list_buffer,
    short result_buffer_length,
    char *result_buffer);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>blob_handle</i>	<i>isc_blob_handle *</i>	Pointer to the Blob
<i>item_list_buffer_length</i>	<i>short</i>	Length of the item-list buffer in which you specify the items for which you want information
<i>item_list_buffer</i>	<i>char *</i>	Pointer to the item-list buffer
<i>result_buffer_length</i>	<i>short</i>	Length of the result buffer into which InterBase returns the requested information
<i>result_buffer</i>	<i>char *</i>	Pointer to the result buffer

Description `isc_blob_info()` returns information about an existing Blob specified by *blob_handle*. The item-list buffer is an unstructured byte vector. An application lists the items about which it wants information in the item-list buffer.

InterBase returns the requested information to the result buffer as a series of clusters of information, one per item requested. Each cluster consists of three parts:

1. A one-byte *item type*. Each is the same as one of the item types in the item-list buffer.
2. A 2-byte number specifying the number of bytes that follow in the remainder of the cluster.
3. A *value*, stored in a variable number of bytes, whose interpretation depends on the item type.

A calling program is responsible for interpreting the contents of the result buffer and for deciphering each cluster as appropriate.

For a list of items that can be requested and returned, see **Chapter 7, “Working with Blob Data.”**

Example The following example retrieves information about the current open Blob:

```
static char blob_items[] = {
    isc_info_blob_max_segment,
    isc_info_blob_num_segments,
    isc_info_blob_ttype};

CHAR blob_info[32];

isc_open_blob2(status_vector, &db, &tr_handle, &blob_handle,
&blob_id, blength, baddr)
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector);
    return(1);
}
isc_blob_info(status_vector, &blob_handle, sizeof(blob_items),
    blob_items, sizeof(blob_info), blob_info);
```

Return Value `isc_blob_info()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_create_blob2()`, `isc_open_blob2()`**

isc_blob_lookup_desc()

Determines the subtype, character set, and segment size of a Blob, given a table name and Blob column name.

```
Syntax  ISC_STATUS isc_blob_lookup_desc(
          ISC_STATUS *status_vector,
          isc_db_handle **db_handle,
          isc_tr_handle **trans_handle,
          unsigned char *table_name,
          unsigned char *column_name,
          ISC_BLOB_DESC *desc,
          unsigned char *global);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle **</i>	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>trans_handle</i>	<i>isc_tr_handle **</i>	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL
<i>table_name</i>	<i>unsigned char *</i>	Name of the table containing the Blob column
<i>column_name</i>	<i>unsigned char *</i>	Name of the Blob column
<i>desc</i>	<i>ISC_BLOB_DESC *</i>	Pointer to the Blob descriptor to which the function returns information
<i>global</i>	<i>unsigned char *</i>	Global column name, returned by this function

Description *isc_blob_lookup_desc()* uses the system tables of a database to determine the subtype, character set, and segment size of a Blob given a table name and Blob column name.

isc_blob_lookup_desc() and three related functions, *isc_blob_default_desc()*, *isc_blob_gen_bpb()*, and *isc_blob_set_desc()* provide dynamic access to Blob information. In particular, you can use these functions to define and access information about Blob data for filtering purposes, such as character set information for text Blob data, and subtype information for text and non-text Blob data.

isc_blob_lookup_desc() stores the requested information about the Blob into the *desc* Blob descriptor structure. The following table describes the *desc* structure:

Parameter	Type	Description
<i>blob_desc_subtype</i>	<i>short</i>	Subtype of the Blob filter
<i>blob_desc_charset</i>	<i>short</i>	Character set being used
<i>blob_desc_segment_size</i>	<i>short</i>	Blob segment size
<i>blob_desc_field_name</i> [32]	<i>char</i>	Array containing the name of the Blob column
<i>blob_desc_relation_name</i> [32]	<i>char</i>	Array containing the name of the table in which the Blob is stored

TABLE 13.18 Blob descriptor fields

Example The following fragment retrieves information into a Blob descriptor:

```
isc_blob_lookup_desc(status, &db_handle, &tr_handle,
    &relation_name, &field_name, desc, &global);
```

Return Value *isc_blob_lookup_desc()* returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code. To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **isc_blob_default_desc(), isc_blob_gen_bpb(), isc_blob_set_desc()**

For more information about Blob descriptors, see **Chapter 7, “Working with Blob Data.”**

isc_blob_set_desc()

Sets the subtype and character set for a Blob.

```
Syntax  ISC_STATUS isc_blob_set_desc(
          ISC_STATUS *status_vector,
          unsigned char *table_name,
          unsigned char *column_name,
          short subtype,
          short charset,
          short segment_size,
          ISC_BLOB_DESC *desc);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>table_name</i>	<i>unsigned char *</i>	Name of the table containing the Blob column
<i>column_name</i>	<i>unsigned char *</i>	Name of the Blob column in the table
<i>subtype</i>	<i>short</i>	Specifies the subtype of the Blob; value are: <ul style="list-style-type: none"> • InterBase-defined subtype values, 0 or 1 (TEXT) • User-defined subtypes, -1 to -32768
<i>charset</i>	<i>short</i>	Specifies the character set for the Blob
<i>segment_size</i>	<i>short</i>	Specifies the segment size for the Blob
<i>desc</i>	<i>ISC_BLOB_DESC *</i>	Pointer to a Blob descriptor to populate

Description *isc_blob_set_desc()* sets the Blob column name, table name, subtype, segment size, and character set for a Blob column to values specified by the application. To set these values to InterBase defaults, use *isc_blob_default_desc()*.

isc_blob_set_desc() and three related functions, *isc_blob_default_desc()*, *isc_blob_gen_bpb()*, and *isc_blob_lookup_desc()* provide dynamic access to Blob data. In particular, you can use these functions to define and access information about Blob data for filtering purposes, such as character set information for text Blob data, and subtype information for text and non-text Blob data.

You can manually set the subtype and character set information (for a TEXT subtype) in a Blob descriptor, by way of a call to *isc_blob_set_desc()*. Pass the subtype, character set, and segment size to the Blob descriptor in your application.

isc_cancel_blob()

isc_blob_set_desc() is useful for setting the contents of the Blob descriptor without querying the system tables for the information. Calls to this function also let an application specify character set and subtype for custom filtering operations.

Note Do not call this function while running against a V3.x database.

Example The following example sets the default values for a tour guide application, including subtype, character set, and segment size:

```
isc_blob_set_desc(status, "TOURISM", "GUIDEBOOK", 1, 2, 80, &desc);
```

Return Value *isc_blob_set_desc()* returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also ***isc_blob_default_desc()***, ***isc_blob_gen_bpb()***, ***isc_blob_lookup_desc()***

For more information about Blob descriptors, see **Chapter 7, “Working with Blob Data.”**

isc_cancel_blob()

Discards a Blob, frees internal storage used by the Blob, and sets the Blob handle to NULL.

Syntax

```
ISC_STATUS isc_cancel_blob(  
    ISC_STATUS *status_vector,  
    isc_blob_handle *blob_handle);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS</i> *	Pointer to the error status vector
<i>blob_handle</i>	<i>isc_blob_handle</i> *	Pointer to the handle for the Blob you want to cancel; sets the handle to zero and returns a successful result even if the handle is NULL

Description InterBase temporarily stores Blob data in the database during create operations. If, for some reason, you do not, or cannot, close a Blob, the storage space remains allocated in the database and InterBase does not set the handle to NULL. Call `isc_cancel_blob()` to release the temporary storage in the database, and to set `blob_handle` to NULL. If you close the Blob in the normal course of your application processing logic, this step is unnecessary as InterBase releases system resources on a call to `isc_close_blob()`.

Note A call to this function does not produce an error when the handle is NULL. Therefore, it is good practice to call `isc_cancel_blob()` before creating or opening a Blob to clean up existing Blob operations.

Example The following fragment cancels any open Blob before creating a new one:

```
isc_cancel_blob(status_vector, &blob_handle);
if (status_vector[0] == 1 && status_vector[1])
{
    /* process error */
    isc_print_status(status_vector);
    return(1);
}
isc_create_blob(status_vector, &DB, &trans, &blob_handle, &blob_id)
```

Return Value `isc_cancel_blob()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_close_blob()`**

isc_cancel_events()

Cancels an application's interest in asynchronous notification of any of a specified group of events.

Syntax `ISC_STATUS isc_cancel_events(
 ISC_STATUS *status_vector,
 isc_db_handle *db_handle,
 ISC_LONG *event_id);`

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle *</i>	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> ; the handle identifies the database for which the event watch is to be canceled. <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>event_id</i>	<i>ISC_LONG *</i>	Pointer to the event or events to cancel; set by a previous call to <i>isc_que_events()</i>

Description *isc_cancel_events()* cancels an application program's asynchronous wait for any of a specified list of events. The events are the ones that were associated with *event_id* as a result of a previous call to *isc_que_events()*.

Example The following call cancels a program's wait for events associated with *event_id*, where *event_id* was previously returned from a call to *isc_que_events()*:

```
isc_cancel_events(status_vector, &database_handle, &event_id);
```

A more complete example is provided in the section on *isc_que_events()*.

Return Value *isc_cancel_events()* returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, "Handling Error Conditions."**

See Also **isc_que_events()**

isc_close_blob()

Closes an open Blob, which involves flushing any remaining segments, releasing system resources associated with Blob update or retrieval, and setting the Blob handle to zero.

Syntax `ISC_STATUS isc_close_blob(
 ISC_STATUS *status_vector,
 isc_blob_handle *blob_handle);`

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>blob_handle</i>	<i>isc_blob_handle *</i>	Pointer to the handle of the Blob to close

Description *isc_close_blob()* is used to store a Blob in the database and clean up after Blob operations. Close any Blob after reading from or writing to it. If, for some reason, your application does not close a Blob, you can lose data. If your application might open a Blob without closing it then you should call *isc_cancel_blob()* to make sure that the application does not try to open a Blob that is already open.

blob_handle is set by a call to *isc_create_blob2()* or to *isc_open_blob2()*.

Example The following example closes a Blob and frees system resources:

```
if (status_vector[1] == isc_segstr_eof)
    isc_close_blob(status_vector, &blob_handle)
```

Return Value *isc_close_blob()* returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also ***isc_cancel_blob()***, ***isc_create_blob2()***, ***isc_open_blob2()***

isc_commit_retaining()

Commits an active transaction and retains the transaction context after a commit.

Syntax `ISC_STATUS isc_commit_retaining(
 ISC_STATUS *status_vector,
 isc_tr_handle *trans_handle);`

Parameter	Type	Description
<i>status_vector</i>	<code>ISC_STATUS *</code>	Pointer to the error status vector
<i>trans_handle</i>	<code>isc_tr_handle *</code>	Pointer to a transaction handle whose value has been set by a previous <code>isc_start_transaction()</code> call; <i>trans_handle</i> returns an error if NULL

Description `isc_commit_retaining()` commits an active transaction and immediately clones itself. This means that the function retains the transaction name, system resources associated with the transaction, and the current state of any open cursors in the transaction. Although the function is actually initiating a new transaction, by assigning the new transaction the active transaction handle it is, in effect, keeping the transaction open across commits. This results in improved performance by allowing an application to minimize the overhead of initiating additional transactions. `isc_commit_retaining()` allows you to commit updates while keeping a cursor open.

You can initiate a rollback within the active transaction but the rollback only affects uncommitted updates. In other words, a rollback is legal, even after the transaction context has been passed to the cloned transaction, but, in that case, the rollback will only affect the updates your application has made to the database since the last commit.

To audit the commits made by your calls to this function, check the first element in the status vector to see if the call was successful. If this element contains a zero, the call was successful.

The transaction ends when you commit or roll back without using the retention feature, with a call to `isc_commit_transaction()` or `isc_rollback_transaction()`.

Examples The following C/C++ code commits a transaction, prints a message, and starts a new transaction with the same handle within the same request:

```
if (!isc_commit_retaining(status, &retained_trans))
{
    fprintf(stderr, "Committed and retained\n");
    isc_print_status(status);
}
```

The following call commits a transaction, prints a confirmation message, starts a new transaction with the same handle within the same request, or, if the commit fails, prints an error message and rolls back.

```
isc_commit_retaining(status, &retained_trans);
if (status[0] == 1 && status[1])
{
    fprintf(stderr, "Error during commit, rolling back.\n");
    rb_status = isc_rollback_transaction(status, &retained_trans);
}
else
{
    fprintf(stderr, "Commit successful.\n");
    tr_count++; /*Increments the number of recycles. */
}
}
```

Return Value `isc_commit_retaining()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_commit_transaction()`**, **`isc_rollback_transaction()`**, **`isc_start_transaction()`**

isc_commit_transaction()

Commits a specified active transaction.

Syntax

```
ISC_STATUS isc_commit_transaction(
    ISC_STATUS *status_vector,
    isc_tr_handle *trans_handle);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>trans_handle</i>	<i>isc_tr_handle *</i>	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL

isc_commit_transaction()

Description `isc_commit_transaction()` closes record streams, frees system resources, and sets the transaction handle to zero for the specified transaction.

When you call this function to execute a commit operation against multiple databases, InterBase first initiates a call to the `isc_prepare_transaction()` function.

`isc_prepare_transaction()` executes the first phase of a two-phase commit. This puts the transaction into limbo and signals your intention to commit, so that InterBase can poll all target databases to verify that they are ready to accept the commit. Also,

`isc_commit_transaction()` writes a Blob message to the RDB\$TRANSACTION_DESCRIPTION column of the RDB\$TRANSACTIONS system table, detailing information required by InterBase to perform a reconnect in case of system failure during the commit process.

The `isc_commit_transaction()` function also performs the second phase of a two-phase commit upon receiving verification that all databases are ready to accept the commit.

Also, `isc_commit_transaction()` cleans up RDB\$TRANSACTIONS.

Example The following call commits a transaction and prints a message:

```
isc_commit_transaction(status, &trans);
if (status[0] == 1 && status[1])
{
    fprintf(stderr, "Error on write\n");
    isc_print_status(status);
}
```

Return Value `isc_commit_transaction()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_commit_retaining()`, `isc_prepare_transaction()`**

isc_create_blob2()

Creates and opens the Blob for write access, and optionally specifies the filters to be used to translate the Blob from one subtype to another.

```
Syntax  ISC_STATUS isc_create_blob2(
          ISC_STATUS *status_vector,
          isc_db_handle *db_handle,
          isc_tr_handle *trans_handle,
          isc_blob_handle *blob_handle,
          ISC_QUAD *blob_id,
          short bpb_length,
          char *bpb_address);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS</i> *	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle</i> *	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>trans_handle</i>	<i>isc_tr_handle</i> *	Pointer to the handle of the transaction in which you want the Blob to be created
<i>blob_handle</i>	<i>isc_blob_handle</i> *	Pointer to the Blob handle
<i>blob_id</i>	<i>ISC_QUAD</i> *	Pointer to the 64-bit system-defined Blob ID, which is stored in a field in the table and points to the first segment of the Blob or to a page of pointers to Blob fragments
<i>bpb_length</i>	<i>short</i>	Length of the Blob parameter buffer (BPB)
<i>bpb_address</i>	<i>char</i> *	Pointer to the BPB

Description *isc_create_blob2()* creates a context for storing a Blob, opens a Blob for write access, and optionally specifies the filters used to translate from one Blob format to another. Subsequent calls to *isc_put_segment()* write data from an application buffer to the Blob. If a Blob filter is used, it is called for each segment written to the Blob. InterBase selects the filter to be used based on the source and target subtypes specified in a previously populated Blob parameter buffer (BPB), pointed to by *bpb_address*.

Note Blob filters are not supported on Netware.

If a Blob filter is not needed or cannot be used, a BPB is not needed; pass 0 for *bpb_length* and NULL for *bpb_address*.

The Blob handle pointed to by *blob_handle* must be zero when *isc_create_blob2()* is called. To reuse *blob_handle*, close the Blob with a call to *isc_close_blob()* to zero out the handle before calling *isc_create_blob2()*.

On success, *isc_create_blob2()* assigns a unique ID to *blob_handle*, and a Blob identifier to *blob_id*. Subsequent API calls require one or both of these to identify the Blob against which they operate.

After a blob is created, data can be written to it by a sequence of calls to *isc_put_segment()*. When finished writing to the Blob, close it with *isc_close_blob()*.

When you create a Blob, it is essentially an “orphan” until you assign its *blob_id* to a particular Blob column of a particular row of a table. You do this, after closing the Blob, by using DSQL to execute either an INSERT statement to insert a new row containing the Blob (and any other columns desired), or an UPDATE statement to replace an existing Blob with the new one.

For more information about BPBs and Blob filters, see **Chapter 7, “Working with Blob Data.”**

Example The following fragment declares a BPB, populates it with filter information, then creates a Blob and passes the BPB:

```
isc_blob_handle blob_handle; /* declare at beginning */
ISC_QUAD blob_id; /* declare at beginning */
char bpb[] = {
    isc_bpb_version1,
    isc_bpb_target_type,
    1,      /* # bytes that follow which specify target subtype */
    1,      /* target subtype (TEXT) */
    isc_bpb_source_type,
    1,      /* # bytes that follow which specify source subtype */
    -4,    /* source subtype*/
};
```

. . .

```
isc_create_blob2(
    status_vector,
    &db_handle,
    &tr_handle,
    &blob_handle, /* to be filled in by this function */
    &blob_id,     /* to be filled in by this function */
```

```

    actual_bpb_length, /* length of BPB data */
    &bpb                /* Blob parameter buffer */
)

```

Return Value `isc_create_blob2()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_blob_gen_bpb()`, `isc_open_blob2()`, `isc_put_segment()`**

isc_create_database()

The `isc_create_database()` method is not currently supported from user applications. It is for internal use only. Use `isc_dsql_execute_immediate()` to create a database with a valid database handle.

isc_database_info()

Reports requested information about a previously attached database.

Syntax

```
ISC_STATUS isc_database_info(
    ISC_STATUS *status_vector,
    isc_db_handle *db_handle,
    short item_list_buffer_length,
    char *item_list_buffer,
    short result_buffer_length,
    char *result_buffer);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle *</i>	Pointer to a database handle set by a previous call to <code>isc_attach_database()</code> <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>item_list_buffer_length</i>	<i>short</i>	Number of bytes in the item-list buffer
<i>item_list_buffer</i>	<i>char *</i>	Address of the item-list buffer
<i>result_buffer_length</i>	<i>short</i>	Number of bytes in the result buffer
<i>result_buffer</i>	<i>char *</i>	Address of the result buffer

Description `isc_database_info()` returns information about an attached database. Typically, `isc_database_info()` is called to:

- Determine how much space is used for page caches. The space is the product of the number of buffers and the page size, which are determined by calling `isc_database_info()` with the `isc_info_num_buffers` and `isc_info_page_size` item-list options.

- Monitor performance. For example, to compare the efficiency of two update strategies, such as updating a sorted or unsorted stream.

The calling program passes its request for information through the item-list buffer supplied by the program, and InterBase returns the information to a program-supplied result buffer.

Example The following program fragment requests the page size and the number of buffers, then examines the result buffer to retrieve the values supplied by the InterBase engine:

```
char db_items[] = {
    isc_info_page_size, isc_info_num_buffers,
    isc_info_end};
char res_buffer[40], *p, item;
int length;
SLONG page_size = 0L, num_buffers = 0L;
ISC_STATUS status_vector[20];

isc_database_info(
    status_vector,
    &handle, /* Set in previous isc_attach_database() call. */
    sizeof(db_items),
    db_items,
    sizeof(res_buffer),
    res_buffer);
if (status_vector[0] == 1 && status_vector[1])
{
    /* An error occurred. */
    isc_print_status(status_vector);
    return(1);
};
/* Extract the values returned in the result buffer. */
for (p = res_buffer; *p != isc_info_end ; )
{
    item = *p++;
    length = isc_vax_integer (p, 2);
    p += 2;
    switch (item)
    {
        case isc_info_page_size:
            page_size = isc_vax_integer (p, length);
            break;
        case isc_info_num_buffers:
```

isc_decode_sql_date()

```
        num_buffers = isc_vax_integer (p, length);
        break;
    default:
        break;
    }
    p += length;
};
```

Return Value `isc_database_info()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_attach_database()`, `isc_detach_database()`**

For more information about requesting database attachment information, see **“Requesting information about an attachment” on page 51.**

isc_decode_sql_date()

Translates a date from InterBase `ISC_DATE` format into the C *struct tm* format.

Syntax

```
void isc_decode_sql_date(
    ISC_DATE *ib_date,
    void *tm_date);
```

Parameter	Type	Description
<i>ib_date</i>	<code>ISC_DATE *</code>	Pointer to a four-byte <code>ISC_DATE</code> structure containing a date in InterBase format
<i>tm_date</i>	<code>void *</code>	Pointer to a C <i>tm</i> structure

Description `isc_decode_sql_date()` translates a date retrieved from a table and stored in an `ISC_DATE` variable, *ib_date*, into a C time structure for program manipulation. Both *ib_date* and *tm_date* must be declared and initialized before use.

Use the `isc_dsql` family of API calls to retrieve InterBase DATE data from a table into the `ISC_DATE` structure prior to translation.

Note In InterBase 6, the DATE datatype is available only in dialect 3. It holds only date information, and does not include time information. In version 6 dialect 1, the TIMESTAMP datatype holds both date and time information and is exactly equivalent to the DATE datatype that was present in earlier versions of InterBase.

Example The following code fragment illustrates declaring time structures and calling `isc_decode_sql_date()` to translate an InterBase date format into a C time format:

```
#include <time.h>
#include <ibase.h>
. . .
struct tm hire_time;
ISC_DATE hire_date;
. . .
/* Retrieve DATE data from a table here. */
. . .
isc_decode_sql_date(&hire_date, &hire_time);
```

Return Value None.

See Also [isc_decode_sql_time\(\)](#), [isc_decode_timestamp\(\)](#), [isc_encode_sql_date\(\)](#)

isc_decode_sql_time()

Translates a time from InterBase `ISC_TIME` format into the C *struct tm* format.

Syntax

```
void isc_decode_sql_time(
    ISC_TIME *ib_time,
    void *tm_date);
```

Parameter	Type	Description
<i>ib_time</i>	<code>ISC_TIME *</code>	Pointer to a four-byte <code>ISC_TIME</code> structure containing a time in InterBase format
<i>tm_date</i>	<code>void *</code>	Pointer to a C <i>struct tm</i> structure

Description `isc_decode_sql_time()` translates a time retrieved from a table and stored in an `ISC_TIME` variable, *ib_time*, into a C time structure for program manipulation. Both *ib_time* and *tm_date* must be declared and initialized before use.

Use the `isc_dsqli` family of API calls to retrieve InterBase TIME data from a table into the `ISC_TIME` structure prior to translation.

Example The following code fragment illustrates declaring time structures and calling `isc_decode_sql_time()` to translate an InterBase date format into a C time format:

```
#include <time.h>
#include <ibase.h>
. . .
struct tm hire_time;
ISC_TIME hire_date;
. . .
/* Retrieve TIME data from a table here. */
. . .
isc_decode_sql_time(&hire_date, &hire_time);
```

Return Value None.

See Also [isc_decode_sql_date\(\)](#), [isc_decode_sql_time\(\)](#), [isc_encode_sql_date\(\)](#)

`isc_decode_timestamp()`

Translates a date and time from InterBase `ISC_TIMESTAMP` format into the C *struct tm* format.

Syntax

```
void isc_decode_timestamp(
    ISC_TIMESTAMP *ib_date,
    void *tm_date);
```

Parameter	Type	Description
<i>ib_timestamp</i>	<code>ISC_TIMESTAMP *</code>	Pointer to an eight-byte <code>ISC_TIMESTAMP</code> structure containing a date and time in InterBase format
<i>tm_date</i>	<code>void *</code>	Pointer to a C <i>struct tm</i> structure

Description `isc_decode_timestamp()` translates a date retrieved from a table and stored in an `ISC_TIMESTAMP` variable, *ib_timestamp*, into a C time structure for program manipulation. Both *ib_timestamp* and *tm_date* must be declared and initialized before use. The `isc_decode_timestamp()` is exactly the same as the `isc_decode_date()` function in versions of InterBase prior to 6.0.

Use the *isc_dsqli* family of API calls to retrieve InterBase `TIMESTAMP` data from a table into the `ISC_TIMESTAMP` structure prior to translation.

Example The following code fragment illustrates declaring time structures and calling `isc_decode_sql_timestamp()` to translate an InterBase date format into a C time format:

```
#include <time.h>
#include <ibase.h>

. . .
struct tm hire_time;
ISC_TIMESTAMP hire_date;

. . .
/* Retrieve TIMESTAMP data from a table here. */
. . .
isc_decode_timestamp(&hire_date, &hire_time);
```

Return Value None.

See Also [isc_decode_sql_date\(\)](#), [isc_decode_sql_time\(\)](#), [isc_encode_sql_date\(\)](#)

isc_delete_user()

Deletes a user record from the password database, `isc4.gdb`.

Note Use of this function is deprecated. It is replaced by a full featured Services API. See [Chapter 12: “Working with Services” on page 199](#) and the reference entry for [“isc_service_start\(\)” on page 380](#).

Syntax

```
ISC_STATUS isc_delete_user(
    ISC_STATUS *status
    USER_SEC_DATA *user_sec_data);
```

Parameter	Type	Description
<i>status vector</i>	<code>ISC_STATUS *</code>	Pointer to the error status vector
<i>user_sec_data</i>	<code>USER_SEC_DATA *</code>	Pointer to a struct that is defined in <code>ibase.h</code>

Description The three security functions, `isc_add_user()`, `isc_delete_user()`, and `isc_modify_user()` mirror functionality that is available in the `gsec` command-line utility. `isc_delete_user()` deletes a record from `isc4.gdb`, InterBase’s password database.

At a minimum, you must provide the user name. If the server is not local, you must provide both a server name and a protocol. Valid choices for the protocol field are `sec_protocol_tcpip`, `sec_protocol_netbeui`, `sec_protocol_spx`, and `sec_protocol_local`.

isc_delete_user()

InterBase reads the settings for the ISC_USER and ISC_PASSWORD environment variables if you do not provide a DBA user name and password.

The definition for the USER_SEC_DATA struct in **ibase.h** is as follows:

```
typedef struct {
    short  sec_flags;      /* which fields are specified */
    int    uid;           /* the user's id */
    int    gid;           /* the user's group id */
    int    protocol;      /* protocol to use for connection */
    char   *server;       /* server to administer */
    char   *user_name;    /* the user's name */
    char   *password;     /* the user's password */
    char   *group_name;   /* the group name */
    char   *first_name;   /* the user's first name */
    char   *middle_name;  /* the user's middle name */
    char   *last_name;    /* the user's last name */
    char   *dba_user_name; /* the dba user name */
    char   *dba_password; /* the dba password */
} USER_SEC_DATA;
```

When you pass this struct to one of the three security functions, you can tell it which fields you have specified by doing a bitwise OR of the following values, which are defined in **ibase.h**:

sec_uid_spec	0x01
sec_gid_spec	0x02
sec_server_spec	0x04
sec_password_spec	0x08
sec_group_name_spec	0x10
sec_first_name_spec	0x20
sec_middle_name_spec	0x40
sec_last_name_spec	0x80
sec_dba_user_name_spec	0x100
sec_dba_password_spec	0x200

No bit values are available for user name and password, since they are required.

The following error messages exist for this function:

Code	Value	Description
<i>isc_username_too_long</i>	335544747	The user name passed in is greater than 31 bytes
<i>isc_password_too_long</i>	335544748	The password passed in is longer than 8 bytes
<i>isc_username_required</i>	335544749	The operation requires a user name
<i>isc_password_required</i>	335544750	The operation requires a password
<i>isc_bad_protocol</i>	335544751	The protocol specified is invalid
<i>isc_dup_username_found</i>	335544752	The user name being added already exists in the security database.
<i>isc_username_not_found</i>	335544753	The user name was not found in the security database
<i>isc_error_adding_sec_record</i>	335544754	An unknown error occurred while adding a user
<i>isc_error_deleting_sec_record</i>	335544755	An unknown error occurred while deleting a user
<i>isc_error_modifying_sec_record</i>	335544756	An unknown error occurred while modifying a user
<i>isc_error_updating_sec_db</i>	335544757	An unknown error occurred while updating the security database

TABLE 13.19 Error messages for user security functions

Example The following example deletes a user (“Socks”) from the password database, using the bitwise OR technique for passing values from the USER_SEC_DATA struct.

```
{
    ISC_STATUS status[20];
    USER_SEC_DATA sec;

    sec.server           = "kennel";
    sec.dba_user_name   = "sysdba";
    sec.dba_password    = "masterkey";
    sec.protocol        = sec_protocol_tcpip;
    sec.user_name       = "socks";
    sec.sec_flags       = sec_server_spec
                        | sec_dba_user_name_spec
                        | sec_dba_password_name_spec;

    isc_delete_user(status, &sec);
}
```

isc_detach_database()

```
/* check status for errors */
if (status[0] == 1 && status[1])
{
    switch (status[1]) {
        case isc_username_too_long:
            printf("Security database cannot accept long user names\n");
            break;
        ...
    }
}
}
```

Return Value `isc_delete_user()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. See the “Description” section for this function for a list of error codes. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_add_user()`**, **`isc_modify_user()`**

isc_detach_database()

Detaches from a database previously connected with `isc_attach_database()`.

Syntax `ISC_STATUS isc_detach_database(
 ISC_STATUS *status_vector,
 isc_db_handle *db_handle);`

Parameter	Type	Description
<i>status_vector</i>	<code>ISC_STATUS *</code>	Pointer to the error status vector
<i>db_handle</i>	<code>isc_db_handle *</code>	Pointer to a database handle set by a previous call to <code>isc_attach_database()</code> <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL

Description `isc_detach_database()` detaches an attached database. Call this function to release system resources when you are done using a database or before re-attaching the database with different attach parameters. `isc_detach_database()` also releases the buffers and structures that control the remote interface on the client and the remote server where the database is stored.

Before calling `isc_detach_database()` commit or roll back transactions affecting the database from which you want to detach.

Example The following conditional statement detaches a database:

```
if (handle)
    isc_detach_database(status_vector, &handle);
```

Assuming that *handle* is valid and identifies an attached database, the specified database is detached when this statement executes.

Return Value `isc_detach_database()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_attach_database()`**

isc_drop_database()

Deletes a currently attached database and all of its supporting files, such as secondary database files, write-ahead log files, and shadow files.

Syntax

```
ISC_STATUS isc_drop_database(
    ISC_STATUS *status_vector,
    isc_db_handle *db_handle);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS</i> *	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle</i> *	Pointer to a database handle set by a previous call to <code>isc_attach_database()</code> ; the handle identifies the database containing the array column <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL

isc_drop_database()

Description `isc_drop_database()` deletes an attached database and all of its supporting files. Call this routine when you no longer have a use for the database (for example, if you moved all the data into another database, or if the database was just temporary and is no longer needed). To succeed, `isc_drop_database()` must be issued when no other processes are attached to the database.

Example The following conditional statement drops a database:

```
if (handle)
    isc_drop_database(status_vector, &handle);
```

Assuming that *handle* is valid and identifies an attached database, the specified database is dropped when this statement executes.

Return Value `isc_drop_database()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_attach_database()`**

isc_dsql_allocate_statement()

Allocates a statement handle for subsequent use with other API dynamic SQL (DSQL) calls.

Syntax `ISC_STATUS isc_dsql_allocate_statement(
 ISC_STATUS *status_vector,
 isc_db_handle *db_handle,
 isc_stmt_handle *stmt_handle);`

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle *</i>	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>stmt_handle</i>	<i>isc_stmt_handle *</i>	Pointer to the statement handle to be allocated by this function; the handle must be NULL when this function is called, or an error is returned in <i>status_vector</i>

Description *isc_dsql_allocate_statement()* allocates a statement handle and returns a pointer to it in *stmt_handle*. This pointer is passed to *isc_dsql_prepare()* to associate the statement handle with a particular DSQL statement for processing.

If a DSQL statement is to be executed multiple times, or if it returns output (other than the results from a stored procedure), *isc_dsql_allocate_statement()* or *isc_dsql_alloc_statement2()* should be called to allocate a statement handle prior to preparing and executing the statement with *isc_dsql_prepare()* and *isc_dsql_execute()*.

Note The function, *isc_dsql_allocate_statement()*, is very similar to the function, *isc_dsql_alloc_statement2()* except that statement handles allocated using *isc_dsql_allocate_statement()* are *not* automatically reset to NULL when the database under which they are allocated is detached. To reset statement handles automatically, use *isc_dsql_alloc_statement2()*.

When you are done processing a statement, the statement handle can be freed with the *isc_dsql_free_statement()* or by calling *isc_detach_database()*.

Example The following program fragment allocates a statement handle for an SQL statement that will access the database referenced by the database handle, *database_handle*:

```
ISC_STATUS status_vector[20];
isc_stmt_handle statement_handle;
```

isc_dsql_allocate_statement()

```
statement_handle = NULL; /* Set handle to NULL before allocating it. */
isc_dsql_allocate_statement(
    status_vector,
    &database_handle, /* Set in previous isc_attach_database() call. */
    &statement_handle);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector); /* Display error message. */
    return(1); /* Return now. */
}

/* Call other functions to associate a particular SQL statement
with the statement handle, and to do other operations necessary to
prepare and execute the DSQL statement. Free the statement handle when
it is no longer needed. */
```

Return Value `isc_dsql_allocate_statement()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to `isc_bad_stmt_handle`, `isc_bad_db_handle`, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_dsql_alloc_statement2()`, `isc_dsql_execute()`, `isc_dsql_free_statement()`, `isc_dsql_prepare()`**

isc_dsql_alloc_statement2()

Allocates a statement handle for subsequent use with other API dynamic SQL (DSQL) calls.

Syntax `ISC_STATUS isc_dsql_alloc_statement2(
 ISC_STATUS *status_vector,
 isc_db_handle *db_handle,
 isc_stmt_handle *stmt_handle);`

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle *</i>	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> ; the handle identifies the database containing the array column <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>stmt_handle</i>	<i>isc_stmt_handle *</i>	Pointer to the statement handle to be allocated by this function; the handle must be NULL when this function is called, or an error is returned in <i>status_vector</i>

Description *isc_dsql_alloc_statement2()* allocates a statement handle and returns a pointer to it in *stmt_handle*. This pointer is passed to *isc_dsql_prepare()* to associate the statement handle with a particular DSQL statement for processing.

If a DSQL statement is to be executed multiple times, or if it returns output (other than the results from a stored procedure), *isc_dsql_alloc_statement2()* or *isc_dsql_allocate_statement()* should be called to allocate a statement handle prior to preparing and executing the statement with *isc_dsql_prepare()* and *isc_dsql_execute()*.

Note The *isc_dsql_allocate_statement2()* function is similar to the *isc_dsql_alloc_statement()* function except that statement handles allocated using *isc_dsql_allocate_statement2()* are automatically reset to NULL when the database under which they are allocated is detached.

Example The following program fragment allocates a statement handle for an SQL statement that will access the database referenced by the database handle, *database_handle*:

```
ISC_STATUS status_vector[20];
isc_stmt_handle statement_handle;

isc_dsql_alloc_statement2(
    status_vector,
```

isc_dsql_alloc_statement2()

```
        &database_handle, /* Set in previous isc_attach_database() call. */
        &statement_handle);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector); /* Display an error message. */
    return(1); /* Return now. */
}
/* Call other functions to associate a particular SQL statement
with the statement handle, and to do other operations necessary to
prepare and execute the DSQL statement. */
;
```

Return Value `isc_dsql_alloc_statement2()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to `isc_bad_stmt_handle`, `isc_bad_db_handle`, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_dsql_allocate_statement()`**, **`isc_dsql_execute()`**, **`isc_dsql_free_statement()`**, **`isc_dsql_prepare()`**

isc_dsql_describe()

Provides information about columns retrieved by the execution of a DSQL SELECT or EXECUTE PROCEDURE statement.

Syntax

```
ISC_STATUS isc_dsql_describe(
    ISC_STATUS *status_vector,
    isc_stmt_handle *stmt_handle,
    unsigned short da_version,
    XSQLDA *xsqlda);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS</i> *	Pointer to the error status vector
<i>stmt_handle</i>	<i>isc_stmt_handle</i> *	Pointer to a statement handle previously allocated with <i>isc_dsql_allocate_statement()</i> or <i>isc_dsql_alloc_statement2()</i> ; the handle returns an error in <i>status_vector</i> if it is NULL
<i>da_version</i>	<i>unsigned short</i>	Indicates the version of the SQL descriptor area passed to the function; set this value to 1
<i>xsqlda</i>	<i>XSQLDA</i> *	Pointer to a previously allocated <i>XSQLDA</i> used for output

Description *isc_dsql_describe()* stores into *xsqlda* a description of the columns that make up the rows returned for a SELECT statement, or a description of the result values returned by an EXECUTE PROCEDURE statement. These statements must have been previously prepared for execution with *isc_dsql_prepare()*, before *isc_dsql_describe()* can be called.

Note Using *isc_dsql_describe()* is not necessary unless a previously issued *isc_dsql_prepare()* function indicates that there is insufficient room in the output *XSQLDA* for the return values of the DSQL statement to be executed.

Example The following program fragment illustrates a sequence of calls which allocates an *XSQLDA*, prepares a statement, checks whether or not the appropriate number of *XSQLVARS* was allocated, and corrects the situation if needed.

```
#include <ibase.h>
ISC_STATUS status_vector[20];
XSQLDA *osqlda;
int n;
char *query = "SELECT * FROM CITIES
              WHERE STATE = 'NY'
              ORDER BY CITY DESCENDING";
```

isc_dsql_describe()

```
osqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(3));
osqlda->version = SQLDA_VERSION1;
osqlda->sqln = 3;

isc_dsql_prepare(
    status_vector,
    &tr_handle, /* Set in previous isc_start_transaction() call. */
    &stmt_handle,
    /* Allocated previously by isc_dsql_allocate_statement()
       or isc_dsql_alloc_statement2() call. */
    0,
    query,
    1,
    osqlda);
if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}

if (osqlda->sqld > osqlda->sqln) /* Need more XSQLVARS. */
{
    n = osqlda->sqld;
    free(osqlda);
    osqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(n));
    osqlda->sqln = n;
    osqlda->version = SQLDA_VERSION1;
    isc_dsql_describe(
        status_vector,
        &stmt_handle,
        1,
        osqlda);
    if (status_vector[0] == 1 && status_vector[1])
    {
        /* Process error. */
        isc_print_status(status_vector);
        return(1);
    }
}
```

Return Value `isc_dsql_describe()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to `isc_bad_stmt_handle`, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_dsql_describe_bind()`, `isc_dsql_execute()`, `isc_dsql_execute2()`, `isc_dsql_prepare()`**

For more information about preparing a DSQL statement with return values, see **“DSQL programming methods” on page 96**. For more information about creating and populating the `XSQLDA`, see **“Understanding the XSQLDA” on page 85**.

`isc_dsql_describe_bind()`

Provides information about dynamic input parameters required by a previously prepared DSQL statement.

Syntax

```
ISC_STATUS isc_dsql_describe_bind(
    ISC_STATUS *status_vector,
    isc_stmt_handle *stmt_handle,
    unsigned short da_version,
    XSQLDA *xsqlda);
```

Parameter	Type	Description
<code>status_vector</code>	<code>ISC_STATUS *</code>	Pointer to the error status vector
<code>stmt_handle</code>	<code>isc_stmt_handle *</code>	Pointer to a statement handle previously allocated with <code>isc_dsql_allocate_statement()</code> or <code>isc_dsql_alloc_statement2()</code> ; the handle returns an error in <code>status_vector</code> if it is NULL
<code>da_version</code>	<code>unsigned short</code>	Indicates the version of the SQL descriptor area passed to the function; set this value to 1
<code>xsqlda</code>	<code>XSQLDA *</code>	Pointer to a previously allocated <code>XSQLDA</code> used for input

Description `isc_dsql_describe_bind()` stores into the input `XSQLDA xsqlda` information about the dynamic input parameters required by a DSQL statement previously prepared with `isc_dsql_prepare()`.

Before an application can execute a statement with input parameters, it must supply values for them in an input *XSQLDA* structure. If you know exactly how many parameters are required, and their datatypes, you can set up the *XSQLDA* directly without calling *isc_dsql_describe_bind()*. But if you need InterBase to analyze the statement and provide information such as the number of parameters and their datatypes, you must call *isc_dsql_describe_bind()* to supply the information.

Example The following program fragment illustrates a sequence of calls that allocates an input *XSQLDA*, prepares a DSQL UPDATE statement, calls the function *isc_dsql_describe_bind()*, checks whether or not the appropriate number of *XSQLVARs* was allocated, and corrects the situation if necessary.

```
#include <ibase.h>
ISC_STATUS status_vector[20];
XSQLDA *isqlda
int n;
char *str = "UPDATE DEPARTMENT SET BUDGET = ?, LOCATION = ?";

isc_dsql_prepare(
    status_vector,
    &tr_handle, /* Set in previous isc_start_transaction() call. */
    &stmt_handle,
    /* Allocated previously by isc_dsql_allocate_statement()
       or isc_dsql_alloc_statement2() call. */
    0,
    str,
    1,
    NULL);
if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}
/* Allocate an input XSQLDA. */
isqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(1));
isqlda->version = SQLDA_VERSION1;
isqlda->sqln = 1;
isc_dsql_describe_bind(
    status_vector,
    &stmt_handle,
    /* Allocated previously by isc_dsql_allocate_statement()
       or isc_dsql_alloc_statement2() call. */
```

```

    1,
    isqlda);
if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}
if (isqlda->sqld > isqlda->sqln) /* Need more XSQLVARs. */
{
    n = isqlda->sqld;
    free(isqlda);
    isqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(n));
    isqlda->sqln = n;
    isqlda->version = SQLDA_VERSION1;
    isc_dsql_describe_bind(
        status_vector,
        &stmt_handle,
        1,
        isqlda);
    if (status_vector[0] == 1 && status_vector[1])
    {
        /* Process error. */
        isc_print_status(status_vector);
        return(1);
    }
}
}

```

Return Value `isc_dsql_describe_bind()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to `isc_bad_stmt_handle`, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_dsql_describe()`, `isc_dsql_execute()`, `isc_dsql_execute2()`, `isc_dsql_prepare()`**

For more information about preparing a DSQL statement with input parameters, see **“DSQL programming methods” on page 96**. For more information about creating and populating the `XSQLDA`, see **“Understanding the XSQLDA” on page 85**.

isc_dsql_execute()

Executes a previously prepared DSQL statement.

Syntax

```
ISC_STATUS isc_dsql_execute(
    ISC_STATUS *status_vector,
    isc_tr_handle *trans_handle,
    isc_stmt_handle *stmt_handle,
    unsigned short da_version,
    XSQLDA *xsqlda);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>trans_handle</i>	<i>isc_tr_handle *</i>	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL
<i>stmt_handle</i>	<i>isc_stmt_handle *</i>	Pointer to a statement handle previously allocated with <i>isc_dsql_allocate_statement()</i> or <i>isc_dsql_alloc_statement2()</i> ; returns an error in <i>status_vector</i> if NULL
<i>da_version</i>	<i>unsigned short</i>	Indicates the version of the extended SQL descriptor area (<i>XSQLDA</i>) passed to the function; set this value to 1
<i>xsqlda</i>	<i>XSQLDA *</i>	Pointer to a previously allocated <i>XSQLDA</i> used for input

Description *isc_dsql_execute()* executes a DSQL statement previously prepared with *isc_dsql_prepare()*. *isc_dsql_execute()* can be used to execute two types of statements:

- Statements that may return more than one row of data.
- Statements that need to be executed more than once.

If a statement to execute has input parameters, then *isc_dsql_execute()* requires an input *XSQLDA* to describe those parameters. It does not provide for an output *XSQLDA*. A call to *isc_dsql_execute()* that executes a SELECT statement results in the creation of a *list* containing all the rows of data that are the result of execution of the statement. To access these rows, call *isc_dsql_fetch()* in a loop. Each call to *isc_dsql_fetch()* fetches the next row from the select-list.

If the statement to be executed requires input parameter values (that is, if it contains parameter markers), these values must be supplied in the input *XSQLDA xsqlda* before calling *isc_dsql_execute()*.

Note To execute a statement repeatedly when it both has input parameters and return values, such as EXECUTE PROCEDURE, use `isc_dsqli_execute2()` which requires both an input and an output *XSQLDA*.

If you only need to execute a statement once, and it does not return any data, call `isc_dsqli_execute_immediate()` instead of `isc_dsqli_prepare()` and `isc_dsqli_execute()`. To execute a statement with both input and output parameters a single time, use `isc_dsqli_exec_immed2()`.

Note CREATE DATABASE and SET TRANSACTION cannot be executed with `isc_dsqli_execute()` or `isc_dsqli_execute2()`. To execute these statements, use `isc_dsqli_execute_immediate()`.

Example The following program fragment illustrates calls to `isc_dsqli_execute()` and `isc_dsqli_fetch()`. It allocates input and output *XSQLDAs*, prepares a SELECT statement, executes it, and fetches and processes each row one-by-one.

```
#include <ibase.h>
ISC_STATUS status_vector[20], fetch_stat;
XSQLDA *isqlda, *osqlda;
XSQLVAR *ivar, *ovar;
char *str = "SELECT CITY, POPULATION FROM CITIES WHERE STATE = ?";
char *state = "CA";
/* Allocate an output XSQLDA osqlda. */
osqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(2));
osqlda->version = SQLDA_VERSION1;
osqlda->sqln = 2;

/* Prepare the statement, including filling in osqlda with information
about the select-list items to be returned by the statement. */
isc_dsqli_prepare(
    status_vector,
    &tr_handle, /* Set in previous isc_start_transaction() call. */
    &stmt_handle,
    /* Allocated previously by isc_dsqli_allocate_statement()
or isc_dsqli_alloc_statement2() call. */
    0,
    str,
    1,
    osqlda);
if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}
```

```
isc_dsql_execute()
```

```
    }

    /* Check to see whether or not the output XSQLDA had enough XSQLVARS
    allocated. If not, correct it -- see isc_dsql_describe(). */

    /* Allocate and fill in the input XSQLDA. This example assumes you know
    how many input parameters there are (1), and all other information
    necessary to supply a value. If this is not true, you will need to call
    isc_dsql_describe_bind(). */
    isqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(1));
    isqlda->version = SQLDA_VERSION1;
    isqlda->sqln = 1;
    isqlda->sqld = 1;
    ivar = isqlda->sqlvar[0];
    ivar->sqltype = SQL_TEXT;
    ivar->sqlllen = sizeof(state);
    ivar->sqldata = state;

    /* Execute the statement. */
    isc_dsql_execute(
        status_vector,
        &tr_handle, /* Set in previous isc_start_transaction() call. */
        &stmt_handle,
        /* Allocated previously by isc_dsql_allocate_statement()
        or isc_dsql_alloc_statement2() call. */
        1,
        isqlda);
    if (status_vector[0] == 1 && status_vector[1])
    {
        /* Process error. */
        isc_print_status(status_vector);
        return(1);
    }

    /* Set up an output XSQLVAR structure to allocate space for each item
    to be returned. */
    for (i=0, ovar = osqlda->sqlvar; i < osqlda->sqld; i++, ovar++)
    {
        dtype = (ovar->sqltype & ~1) /* Drop NULL bit for now. */
        switch(dtype)
        {
            case SQL_TEXT:
```

```

        ovar->sqldata = (char *)malloc(sizeof(char) * ovar->sqlen);
        break;
    case SQL_LONG:
        ovar->sqldata = (char *)malloc(sizeof(long));
        /* Process remaining types. */
        . . .
    }
    if (ovar->sqltype & 1)
    {
        /* Assign a variable to hold NULL status. */
        ovar->sqlind = (short *)malloc(sizeof(short));
    }
} /* end of for loop */

/* Fetch and process the rows in the select list one by one. */
while ((fetch_stat = isc_dsql_fetch(
    status_vector,
    &stmt_handle,
    1,
    osqlda)) == 0)
{
    for (i=0; i < osqlda->sqld; i++)
    {
        /* Call a function you've written to process each returned
        select-list item. */
        process_column(osqlda->sqlvar[i]);
    }
}

```

Return Value `isc_dsql_execute()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to *isc_bad_stmt_handle*, *isc_bad_trans_handle*, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_dsql_describe_bind()`, `isc_dsql_exec_immed2()`, `isc_dsql_execute_immediate()`, `isc_dsql_execute2()`, `isc_dsql_fetch()`, `isc_dsql_prepare()`**

For more information about creating and populating the *XSQLDA*, see **“Understanding the XSQLDA” on page 85.**

isc_dsql_execute2()

Executes a previously prepared DSQL statement.

Syntax

```
ISC_STATUS isc_dsql_execute2(
    ISC_STATUS *status_vector,
    isc_tr_handle *trans_handle,
    isc_stmt_handle *stmt_handle,
    unsigned short da_version,
    XSQLDA *in_xsqlda,
    XSQLDA *out_xsqlda);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS</i> *	Pointer to the error status vector
<i>trans_handle</i>	<i>isc_tr_handle</i> *	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL
<i>stmt_handle</i>	<i>isc_stmt_handle</i> *	Pointer to a statement handle previously allocated with <i>isc_dsql_allocate_statement()</i> or <i>isc_dsql_alloc_statement2()</i> ; the handle returns an error in <i>status_vector</i> if it is NULL
<i>da_version</i>	<i>unsigned short</i>	Indicates the version of the extended SQL descriptor area (<i>XSQLDA</i>) passed to the function; set this value to 1
<i>in_xsqlda</i>	<i>XSQLDA</i> *	Pointer to an optional, previously allocated <i>XSQLDA</i> used for input; if input parameters are not supplied, set this value to NULL
<i>out_xsqlda</i>	<i>XSQLDA</i> *	Pointer to an optional, previously allocated <i>XSQLDA</i> used for results of statement execution; if not required, set this value to NULL

Description *isc_dsql_execute2()* executes a previously prepared DSQL statement that has input parameters and returns results, such as EXECUTE PROCEDURE and SELECT.

If the statement to execute requires input parameter values (that is, if it contains parameter markers), these values must be supplied in the input *XSQLDA*, *in_xsqlda* before calling *isc_dsql_execute2()*.

If the statement to execute returns values, they are placed in the specified output *XSQLDA*, *out_xsqlda*. If a NULL value is supplied for the output *XSQLDA* and the statement returns values, they are stored in a *result set*. To access the returned data, use *isc_dsql_fetch()* in a loop.

Tip If you just want to execute once a statement returning just one group of data, call *isc_dsql_exec_immed2()* instead of *isc_dsql_prepare()* and *isc_dsql_execute2()*.

To execute a statement that does not return any data a single time, call *isc_dsql_execute_immediate()* instead of *isc_dsql_prepare()* and *isc_dsql_execute2()*.

Note CREATE DATABASE and SET TRANSACTION cannot be executed with *isc_dsql_execute()* or *isc_dsql_execute2()*. To execute these statements, use *isc_dsql_execute_immediate()*.

Example The following program fragment illustrates a sequence of calls that allocates an input *XSQLDA* and loads values into it, allocates an output *XSQLDA*, prepares an EXECUTE PROCEDURE statement, allocates space in the output *XSQLDA* for each column returned for each row retrieved by the call, and executes the prepared statement, placing return values in the output *XSQLDA*.

```
#include <ibase.h>
ISC_STATUS status_vector[20];
XSQLDA *isqlda, *osqlda;
XSQLVAR *ivar, *ovar;
short null_flag;
char *str = "EXECUTE PROCEDURE P1";
char *state = "CA";
/* Allocate an output XSQLDA osqlda. This example assumes you know that
P1 will return one value. */
osqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(1));
osqlda->version = SQLDA_VERSION1;
osqlda->sqln = 1;

/* Prepare the statement, including filling in osqlda with information
about the item to be returned by the statement (procedure). */
isc_dsql_prepare(
    status_vector,
    &tr_handle, /* Set in previous isc_start_transaction() call. */
    &stmt_handle,
    /* Allocated previously by isc_dsql_allocate_statement()
or isc_dsql_alloc_statement2() call. */
    0,
    str,
    1,
```

```

    osqlda);
if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}
/* Set up the output XSQLVAR structure to allocate space for the return
value. Again, this example assumes you know that P1 returns just one
value. For an example of what to do if you're not sure, see
isc_dsql_describe(). For an example of setting up an output XSQLVAR
structure to allocate space for multiple return items, see the
isc_dsql_execute() example program. */
ovar = osqlda->sqlvar[0];
dtype = (ovar->sqltype & ~1); /* Drop NULL bit for now. */
switch(dtype)
{
    case SQL_TEXT:
        ovar->sqldata = (char *)malloc(sizeof(char) * ovar->sqlen);
        break;
    case SQL_LONG:
        ovar->sqldata = (char *)malloc(sizeof(long));
        /* Process remaining types. */
        . . .
}
if (ovar->sqltype & 1)
{
    /* Assign a variable to hold NULL status. */
    ovar->sqlind = &null_flag;
}
/* Allocate and fill in the input XSQLDA. This example assumes you know
how many input parameters there are (1), and all other information
necessary to supply a value. If this is not true, you will need to call
isc_dsql_describe_bind(). */
isqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(1));
isqlda->version = SQLDA_VERSION1;
isqlda->sqln = 1;
isqlda->sqld = 1;
ivar = isqlda->sqlvar[0];
ivar->sqltype = SQL_TEXT;
ivar->sqlen = sizeof(state);
ivar->sqldata = state;

```

```

/* Execute the statement. */
isc_dsqli_execute2(
    status_vector,
    &tr_handle, /* Set in previous isc_start_transaction() call. */
    &stmt_handle,
    /* Allocated previously by isc_dsqli_allocate_statement()
       or isc_dsqli_alloc_statement2() call. */
    1,
    isqllda,
    osqllda);
if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}
/* Now process the value returned in osqllda->sqlvar[0]. */
. . .

```

Return Value `isc_dsqli_execute2()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to `isc_bad_stmt_handle`, `isc_bad_trans_handle`, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_dsqli_exec_immed2()`, `isc_dsqli_execute_immediate()`, `isc_dsqli_execute()`, `isc_dsqli_fetch()`, `isc_dsqli_prepare()`**

For more information about creating and populating the `XSQLDA`, see **“Understanding the XSQLDA” on page 85.**

isc_dsqli_execute_immediate()

Prepares and executes just once a DSQL statement that does not return data. There is a special case of *isc_dsqli_execute_immediate()* for creating databases.

Syntax

```
ISC_STATUS isc_dsqli_execute_immediate(
    ISC_STATUS *status_vector,
    isc_db_handle *db_handle,
    isc_tr_handle *trans_handle,
    unsigned short length,
    char *statement,
    unsigned short dialect,
    XSQLDA *xsqlda);
```

Note In the special case where the statement is CREATE DATABASE, there is no transaction, so *db_handle* and *trans_handle* must be pointers to handles whose value is NULL. When *isc_dsqli_execute_immediate()* returns, *db_handle* is a valid handle, just as though you had made a call to *isc_attach_database()*.

Parameter	Type	Description
<i>status_vector</i>	ISC_STATUS *	Pointer to the error status vector
<i>db_handle</i>	isc_db_handle *	<ul style="list-style-type: none"> If <i>statement</i> is not CREATE DATABASE, this is a pointer to a database handle set by a previous call to <i>isc_attach_database()</i>; <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL If <i>statement</i> is CREATE DATABASE, this must point to a database handle whose value is NULL
<i>trans_handle</i>	isc_tr_handle *	<ul style="list-style-type: none"> If <i>statement</i> is not CREATE DATABASE, this is a pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL If <i>statement</i> is CREATE DATABASE or SET TRANSACTION, this must point to a transaction handle whose value is NULL
<i>length</i>	unsigned short	Length of the DSQL statement in bytes; set to 0 in C programs to indicate a null-terminated string

Parameter	Type	Description
<i>statement</i>	<i>char *</i>	DSQL string to be executed
<i>dialect</i>	<i>unsigned short</i>	<ul style="list-style-type: none"> • Indicates the SQL dialect of <i>statement</i> • Must be less than or equal to the SQL dialect of the client
<i>xsqlda</i>	<i>XSQLDA *</i>	Pointer to an optional, previously allocated <i>XSQLDA</i> used for input; if you don't supply input parameters, set this value to NULL

Description *isc_dsql_execute_immediate()* prepares the DSQL statement specified in *statement*, executes it once, and discards it. The statement must not be one that returns data (that is, it must not be a SELECT or EXECUTE PROCEDURE statement).

If *statement* requires input parameter values (that is, if it contains parameter markers), these values must be supplied in the input *XSQLDA*, *xsqlda*.

To create a database using *isc_dsql_execute_immediate()*, supply a CREATE DATABASE statement and have *db_handle* and *trans_handle* point to handles with a NULL value.

Tip If *statement* returns data, or if it needs to be executed more than once, use *isc_dsql_prepare()* and *isc_dsql_execute()* (or *isc_dsql_execute2()*) instead of *isc_dsql_execute_immediate()*.

Note You *must* call *isc_dsql_execute_immediate()* rather than *isc_dsql_prepare()* and *isc_dsql_execute()* for CREATE DATABASE or SET TRANSACTION. To start a transaction, you also have the option of using *isc_start_transaction()*.

Examples The following program fragment calls *isc_dsql_execute_immediate()* to perform an insert:

```
#include <ibase.h>
ISC_STATUS status_vector[20];
char *insert_stmt =
    "INSERT INTO CUSTOMER(CUSTNAME, BAL, CUSTNO)
    VALUES('John Smith', 299.0, 5050)";

isc_dsql_execute_immediate(
    status_vector,
    &database_handle, /* Set in previous isc_attach_database() call. */
    &tr_handle, /* Set in previous isc_start_transaction() call. */
    0,
    insert_stmt,
    1,
    NULL);
```

isc_dsqli_execute_immediate()

```
if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}
```

The following C/C++ code fragment uses *isc_dsqli_execute_immediate()* to create a database and return a handle to the new database:

```
#include <ibase.h>
ISC_STATUS status_vector[20];
char *statement =
    "CREATE DATABASE 'C:/INVENTORY.GDB' PAGE_SIZE 4096 \
    USER 'SYSDBA' PASSWORD 'masterkey' ";
isc_db_handle db_handle = NULL;
isc_tr_handle dummy_handle = NULL;

isc_dsqli_execute_immediate(
    status_vector,
    &db_handle,
    &dummy_handle,
    0,
    statement,
    1,
    NULL);
if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}
```

Return Value *isc_dsqli_execute_immediate()* returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to *isc_bad_db_handle*, *isc_bad_trans_handle*, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also ***isc_dsqli_exec_immed2()***, ***isc_dsqli_execute()***, ***isc_dsqli_prepare()***

For more information about creating and populating the *XSQLDA*, see **“Understanding the XSQLDA” on page 85**.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

isc_dsql_exec_immed2()

Prepares and executes just once, a DSQL statement that returns no more than one row of data.

```
Syntax  ISC_STATUS isc_dsql_exec_immed2(
          ISC_STATUS *status_vector,
          isc_db_handle *db_handle,
          isc_tr_handle *trans_handle,
          unsigned short length,
          char *statement,
          unsigned short dialect,
          XSQLDA *in_xsqlda,
          XSQLDA *out_xsqlda);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS</i> *	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle</i> *	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>trans_handle</i>	<i>isc_tr_handle</i> *	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL
<i>length</i>	<i>unsigned short</i>	Length of the DSQL statement, in bytes; set to 0 in C programs to indicate a null-terminated string
<i>statement</i>	<i>char</i> *	DSQL string to be executed

Parameter	Type	Description
<i>dialect</i>	<i>unsigned short</i>	<ul style="list-style-type: none"> Indicates the SQL dialect of statement Must be less than or equal to the SQL dialect of the client
<i>in_xsqlda</i>	<i>XSQLDA *</i>	Pointer to an optional, previously allocated <i>XSQLDA</i> used for input; if input parameters are not supplied, set this value to NULL
<i>out_xsqlda</i>	<i>XSQLDA *</i>	Pointer to an optional, previously allocated <i>XSQLDA</i> used for results of statement execution. If not required, set this value to NULL

Description *isc_dsql_exec_immed2()* prepares the DSQL statement specified in *statement*, executes it once, and discards it. *statement* can return a single set of values (i.e, it can be an EXECUTE PROCEDURE or singleton SELECT) in the output *XSQLDA*.

If *statement* requires input parameter values (that is, if it contains parameter markers), these values must be supplied in the input *XSQLDA*, *in_xsqlda*.

For statements that return multiple rows of data, use *isc_dsql_prepare()*, *isc_dsql_execute2()*, and *isc_dsql_fetch()*.

Example The following program fragment calls *isc_dsql_exec_immed2()*:

```
ISC_STATUS status_vector[20];
XSQLDA *in_xsqlda, *out_xsqlda;
char *execute_p1 = "EXECUTE PROCEDURE P1 ?";
/* Set up input and output XSQLDA structures here. */
. . .
isc_dsql_exec_immed2(
    status_vector,
    &database_handle, /* Set in previous isc_attach_database() call. */
    &tr_handle, /* Set in previous isc_start_transaction() call. */
    0,
    execute_p1,
    1,
    in_xsqlda,
    out_xsqlda);
if (status_vector[0] == 1 && status_vector[1]) {
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}
```

Return Value `isc_dsql_exec_immed2()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to `isc_bad_db_handle`, `isc_bad_trans_handle`, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_dsql_execute2()`, `isc_dsql_prepare()`**

For more information about creating and populating the `XSQLDA`, see **“Understanding the `XSQLDA`” on page 85.**

`isc_dsql_fetch()`

Retrieves data returned by a previously prepared and executed DSQL statement.

Syntax

```
ISC_STATUS isc_dsql_fetch(
    ISC_STATUS *status_vector,
    isc_stmt_handle *stmt_handle,
    unsigned short da_version,
    XSQLDA *xsqlda);
```

Parameter	Type	Description
<code>status_vector</code>	<code>ISC_STATUS *</code>	Pointer to the error status vector
<code>stmt_handle</code>	<code>isc_stmt_handle *</code>	Pointer to a statement handle previously allocated with <code>isc_dsql_allocate_statement()</code> or <code>isc_dsql_alloc_statement2()</code> ; the handle returns an error in <code>status_vector</code> if it is NULL
<code>da_version</code>	<code>unsigned short</code>	Indicates the version of the extended SQL descriptor area (<code>XSQLDA</code>) passed to the function; set this value to 1
<code>xsqlda</code>	<code>XSQLDA *</code>	Pointer to an optional, previously allocated <code>XSQLDA</code> used for results of statement execution

Description `isc_dsql_fetch()` retrieves one row of data into `xsqlda` each time it is called. It is used in a loop to retrieve and process each row of data for statements that return multiple rows in a cursor.

A cursor is a one-way pointer into the ordered set of rows retrieved by a statement. A cursor is only needed to process positioned UPDATE and DELETE statements made against the rows retrieved by *isc_dsqli_fetch()* for SELECT statements that specify an optional FOR UPDATE OF clause.

It is up to the application to provide the loop construct for fetching the data.

Before calling *isc_dsqli_fetch()*, a statement must be prepared with *isc_dsqli_prepare()*, and executed with *isc_dsqli_execute()* (or *isc_dsqli_execute2()* with a NULL output *xsqllda* argument). Statement execution produces a result set containing the data returned. Each call to *isc_dsqli_fetch()* retrieves the next available row of data from the result set into *xsqllda*.

Example The following program fragment illustrates a sequence of calls that allocates an output *XSQLDA*, prepares a statement for execution, allocates an *XSQLVAR* structure in the *XSQLDA* for each column of data to be retrieved, executes the statement, producing a select list of returned data, then fetches and processes each row in a loop:

```
#include <ibase.h>
#define LASTLEN 20
#define FIRSTLEN 15
#define EXTLEN 4
typedef struct vary {
    short vary_length;
    char vary_string[1];
} VARY;
ISC_STATUS status_vector[20], retcode;
long SQLCODE;
XSQLDA *osqlda;
XSQLVAR *ovar;
short flag0, flag1, flag2;
char *str =
    "SELECT last_name, first_name, phone_ext FROM phone_list
    WHERE location = \"Monterey\" ORDER BY last_name, first_name";
char last_name[LASTLEN + 2];
char first_name[FIRSTLEN + 2];
char phone_ext[EXTLEN + 2];
VARY *vary;
/* Allocate an output XSQLDA osqlda. */
osqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(3));
osqlda->version = SQLDA_VERSION1;
osqlda->sqln = 3;
/* Prepare the statement. */
isc_dsqli_prepare(
```

```

    status_vector,
    &tr_handle, /* Set in previous isc_start_transaction() call. */
    &stmt_handle,
    /* Allocated previously by isc_dsqli_allocate_statement()
       or isc_dsqli_alloc_statement2() call. */
    0,
    str,
    1,
    osqlda);
if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}
/* Set up an output XSQLVAR structure to allocate space for each item
to be returned. */
osqlda->sqlvar[0].sqldata = last_name;
osqlda->sqlvar[0].sqltype = SQL_VARYING + 1;
osqlda->sqlvar[0].sqlind = &flag0;
osqlda->sqlvar[1].sqldata = first_name;
osqlda->sqlvar[1].sqltype = SQL_VARYING + 1;
osqlda->sqlvar[1].sqlind = &flag1;
osqlda->sqlvar[2].sqldata = phone_ext;
osqlda->sqlvar[2].sqltype = SQL_VARYING + 1;
osqlda->sqlvar[2].sqlind = &flag2;
/* Execute the statement. */
isc_dsqli_execute(
    status_vector,
    &tr_handle, /* Set in previous isc_start_transaction() call. */
    &stmt_handle,
    /* Allocated previously by isc_dsqli_allocate_statement()
       or isc_dsqli_alloc_statement2() call. */
    1,
    NULL);
if (status_vector[0] == 1 && status_vector[1])
{
    /* Process error. */
    isc_print_status(status_vector);
    return(1);
}

```

isc_dsql_fetch()

```
printf("\n%-20s %-15s %-10s\n\n", "LAST NAME", "FIRST NAME",
"EXTENSION");
/* Fetch and print the records in the select list one by one. */
while ((retcode = isc_dsql_fetch(
    status_vector,
    &stmt_handle,
    1,
    osqllda)) == 0)
{
    vary = (VARY *)last_name;
    printf("%-20.*s ", vary->vary_length, vary->vary_string);
    vary = (VARY *)first_name;
    printf("%-15.*s ", vary->vary_length, vary->vary_string);
    vary = (VARY *)phone_ext;
    printf("%-4.*s ", vary->vary_length, vary->vary_string);
}
if (retcode != 100L)
{
    SQLCODE = isc_sqlcode(status_vector);
    isc_print_sqlerror(SQLCODE, status_vector);
    return(1);
}
```

Return Value `isc_dsql_fetch()` returns the second element of the status vector. Zero indicates success. The value 100 indicates that no more rows remain to be retrieved. Any other nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to `isc_bad_stmt_handle`, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_dsql_execute()`, `isc_dsql_execute2()`, `isc_dsql_prepare()`**

isc_dsql_free_statement()

Frees a statement handle and all resources allocated for it, or closes a cursor associated with the statement referenced by a statement handle.

Syntax `ISC_STATUS isc_dsql_free_statement(
 ISC_STATUS *status_vector,
 isc_stmt_handle *stmt_handle,
 unsigned short option);`

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>stmt_handle</i>	<i>isc_stmt_handle *</i>	Pointer to a statement handle previously allocated with <i>isc_dsql_allocate_statement()</i> or <i>isc_dsql_alloc_statement2()</i> ; the handle returns an error in <i>status_vector</i> if it is NULL
<i>option</i>	<i>unsigned short</i>	Either <i>DSQL_close</i> or <i>DSQL_drop</i>

Description *isc_dsql_free_statement()* either frees a statement handle and all resources allocated for it (*option = DSQL_drop*), or closes a cursor associated with the statement (*option = DSQL_close*).

Note *isc_dsql_free_statement()* does nothing if it is called with an *option* value other than *DSQL_drop* or *DSQL_close*.

► *DSQL_close*

Call *isc_dsql_free_statement()* with the *DSQL_close* option to close a cursor after it is no longer needed, that is, after fetching and processing all the rows resulting from the execution of a query. A cursor need only be closed in this manner if it was previously opened and associated with *stmt_handle* by *isc_dsql_set_cursor_name()*.

DSQL_close closes a cursor, but the statement it was associated with remains available for further execution.

If you have used a cursor to perform updates or deletes on all the rows returned from the execution of a query, and you want to perform other update or delete operations on rows resulting from execution of the same statement again (possibly with different input parameters), follow these steps:

1. Close the cursor with *isc_dsql_free_statement()*.
2. Re-open it with *isc_dsql_set_cursor_name()*.

3. If desired, change the input parameters to be passed to the statement.
4. Re-execute the statement to retrieve a new select list.
5. Retrieve rows in a loop with *isc_dsql_fetch()* and process them again with *isc_dsql_execute_immediate()*.

► *DSQL_drop*

Statement handles allocated with *isc_dsql_allocate_statement()* must be released when no longer needed by calling *isc_dsql_free_statement()* with the *DSQL_drop* option. This option frees all resources associated with the statement handle, and closes any open cursors associated with the statement handle.

Example The following program fragment shows examples of the two types of *isc_dsql_free_statement()* calls. It assumes that *stmt_handle1* and *stmt_handle2* are statement handles, each of which was previously allocated with either *isc_dsql_allocate_statement()* or *isc_dsql_alloc_statement2()*. A cursor is also assumed to have been associated with the statement referenced by *stmt_handle1*.

```
#include <ibase.h>
ISC_STATUS status_vector[20];
. . .
/* Free the cursor associated with stmt_handle1. */
isc_dsql_free_statement(
    status_vector,
    &stmt_handle1,
    DSQL_close);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector);
    return(1);
}
/* Free stmt_handle2. */
isc_dsql_free_statement(
    status_vector,
    &stmt_handle2,
    DSQL_drop);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector);
    return(1);
}
```

Return Value `isc_dsql_free_statement()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to `isc_bad_stmt_handle`, or another InterBase error code. To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_dsql_allocate_statement()`**, **`isc_dsql_alloc_statement2()`**,
`isc_dsql_set_cursor_name()`

`isc_dsql_prepare()`

Prepares a DSQL statement for repeated execution.

Syntax

```
ISC_STATUS isc_dsql_prepare(
    ISC_STATUS *status_vector,
    isc_tr_handle *trans_handle,
    isc_stmt_handle *stmt_handle,
    unsigned short length,
    char *statement,
    unsigned short dialect,
    XSQLDA *xsqllda);
```

Parameter	Type	Description
<code>status_vector</code>	<code>ISC_STATUS *</code>	Pointer to the error status vector
<code>trans_handle</code>	<code>isc_tr_handle *</code>	Pointer to a transaction handle whose value has been set by a previous <code>isc_start_transaction()</code> call; <code>trans_handle</code> returns an error if NULL
<code>stmt_handle</code>	<code>isc_stmt_handle *</code>	Pointer to a statement handle previously allocated with <code>isc_dsql_allocate_statement()</code> or <code>isc_dsql_alloc_statement2()</code> ; the handle returns an error in <code>status_vector</code> if it is NULL
<code>length</code>	<code>unsigned short</code>	Length of the DSQL statement, in bytes; set to 0 in C programs to indicate a null-terminated string

Parameter	Type	Description
<i>statement</i>	<i>char *</i>	DSQL string to be executed
<i>dialect</i>	<i>unsigned short</i>	<ul style="list-style-type: none"> Indicates the SQL dialect of statement Must be less than or equal to the SQL dialect of the client
<i>xsqlda</i>	<i>XSQLDA *</i>	Pointer to an optional, previously allocated XSQLDA used for results of statement execution

Description *isc_dsql_prepare()* readies the DSQL statement specified in *statement* for repeated execution by checking it for syntax errors and parsing it into a format that can be efficiently executed. All SELECT statements must be prepared with *isc_dsql_prepare()*.

After a statement is prepared, it is available for execution as many times as necessary during the current session. Preparing a statement for repeated execution is more efficient than using *isc_dsql_execute_immediate()* or *isc_dsql_exec_immed2()* over and over again to prepare and execute a statement.

If a statement to be prepared does not return data, set the output *XSQLDA* to NULL. Otherwise, the output *XSQLDA* must be allocated prior to calling *isc_dsql_prepare()*. Allocate the *XSQLDA* using the macro, *XSQLDA_LENGTH*, defined in **ibase.h**, as follows:

```
xsqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(n));
```

XSQLDA_LENGTH calculates the number of bytes required when *n* result columns will be returned by the statement, and allocates the appropriate amount of storage.

After allocating the *XSQLDA xsqlda*, set *xsqlda->version* to *SQLDA_VERSION1*, and set *xsqlda_sqln* to indicate the number of *XSQLVAR* structures allocated.

When *isc_dsql_prepare()* is called, it fills in the other fields of the *XSQLDA* and all the *XSQLVARs* with information such as the datatype, length, and name of the corresponding select-list items in the statement. It fills in *xsqlda->sqld* with the actual number of select-list items returned. If *xsqlda->sqld* is greater than *xsqlda->sqln*, then enough room is not allocated, and the *XSQLDA* must be resized by following these steps:

1. Record the current value of the *xsqlda->sqld*.
2. Free the storage previously allocated for *xsqlda*.
3. Reallocate storage for *xsqlda*, this time specifying the correct number (from step 1) in the argument to *XSQLDA_LENGTH*.
4. Reset *xsqlda->sqld* and *xsqlda->version*.
5. Execute *isc_dsql_describe()* to fill in the *xsqlda* fields.

Note If the prepared statement requires input parameter values, then an input *XSQLDA* will need to be allocated and filled in with appropriate values prior to calling *isc_dsqli_execute()* or *isc_dsqli_execute2()*. You can either allocate and directly fill in all the fields of the input *XSQLDA*, or you can allocate it, call *isc_dsqli_describe_bind()* to get information regarding the number and types of parameters required, then fill in appropriate values.

Example The following program fragment illustrates the allocation of the output *XSQLDA*, and a call to *isc_dsqli_prepare()*:

```
#include <ibase.h>
ISC_STATUS status_vector[20];
XSQLDA *osqlda;
char *query = "SELECT CITY, STATE, POPULATION FROM CITIES \
              WHERE STATE = \"NY\" ORDER BY CITY DESCENDING";

osqlda = (XSQLDA *)malloc(XSQLDA_LENGTH(3));
osqlda->version = SQLDA_VERSION1;
osqlda->sqln = 3;

isc_dsqli_prepare(
    status_vector,
    &tr_handle, /* Set in previous isc_start_transaction() call. */
    &stmt_handle,
    /* Allocated previously by isc_dsqli_allocate_statement()
       or isc_dsqli_alloc_statement2() call. */
    0,
    query,
    1,
    osqlda);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector);
    return(1);
}
```

More complete examples showing the subsequent execution and fetching of result data are provided in the example programs for *isc_dsqli_execute()*, *isc_dsqli_execute2()*, and *isc_dsqli_fetch()*.

isc_dsqli_set_cursor_name()

Return Value `isc_dsqli_prepare()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to `isc_bad_stmt_handle`, `isc_bad_trans_handle`, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_dsqli_describe()`, `isc_dsqli_describe_bind()`, `isc_dsqli_execute()`, `isc_dsqli_execute2()`, `isc_dsqli_fetch()`**

For more information about creating and populating the `XSQLDA`, see **“Understanding the XSQLDA” on page 85 of Chapter 6, “Working with Dynamic SQL.”**

`isc_dsqli_set_cursor_name()`

Defines a cursor name and associates it with a DSQL statement.

Syntax

```
ISC_STATUS isc_dsqli_set_cursor_name(  
    ISC_STATUS *status_vector,  
    isc_stmt_handle *stmt_handle,  
    char *cursor_name,  
    unsigned short type);
```

Parameter	Type	Description
<code>status_vector</code>	<code>ISC_STATUS *</code>	Pointer to the error status vector
<code>stmt_handle</code>	<code>isc_stmt_handle *</code>	Pointer to a statement handle previously allocated with <code>isc_dsqli_allocate_statement()</code> or <code>isc_dsqli_alloc_statement2()</code> ; the handle returns an error in <code>status_vector</code> if it is NULL
<code>cursor_name</code>	<code>char *</code>	String name of a cursor
<code>type</code>	<code>unsigned short</code>	Reserved for future use; set to NULL

Description `isc_dsqli_set_cursor_name()` defines a cursor name and associates it with a DSQL statement handle for a statement that returns multiple rows of data (for example, SELECT), effectively opening the cursor for access.

A cursor is a one-way pointer into the ordered set of rows retrieved by a statement. A cursor is only needed to process positioned UPDATE and DELETE statements made against the rows retrieved by *isc_dsql_fetch()* for SELECT statements that specify an optional FOR UPDATE OF clause.

Note In UPDATE or DELETE statements, the cursor name cannot be supplied as a parameter marker (?).

When a cursor is no longer needed, close it with the *DSQL_close* option of *isc_dsql_free_statement()*.

Example The following pseudo-code illustrates the calling sequence necessary to execute an UPDATE or DELETE with the WHERE CURRENT OF clause using a cursor name established and opened with *isc_dsql_set_cursor_name()*:

```
#include <ibase.h>
ISC_STATUS status_vector[20], fetch_stat;
isc_stmt_handle st_handle = NULL;
char *cursor = "S";

/* Allocate the statement handle st_handle. */
isc_dsql_allocate_statement(
    status_vector,
    &db, /* Database handle set by isc_attach_database() call. */
    &st_handle);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector);
    return(1);
}
/* Set up an output XSQLDA osqlda here. */
/* Call isc_dsql_prepare() to prepare the SELECT statement. */
/* Set up an input XSQLDA, if needed, for the SELECT statement. */
/* Call isc_dsql_execute() to execute the SELECT statement. */
/* Set up an input XSQLDA (if needed) for the UPDATE or DELETE
statement. */
/* Declare the cursor name, and associate it with st_handle. */
isc_dsql_set_cursor_name(
    status_vector,
    &st_handle,
    cursor, 0);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector);
}
```

isc_dsql_set_cursor_name()

```
        return(1);
    }
    /* Fetch rows one by one, with the cursor pointing to each row as it
    is fetched, and execute an UPDATE or DELETE statement to update or
    delete the row pointed to by the cursor. */
    while ((fetch_stat = isc_dsql_fetch(
        status_vector, &st_handle, 1, osqlda)) == 0)
    {
        . . .
        /* Update or delete the current row by executing an "UPDATE ...
        WHERE CURRENT OF S" or "DELETE ... WHERE CURRENT OF S"
        statement, where "S" is the name of the cursor declared in
        isc_dsql_set_cursor_name(). */
    }
}
```

Return Value `isc_dsql_set_cursor_name()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to *isc_bad_stmt_handle*, or another InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_dsql_fetch()`, `isc_dsql_free_statement()`**

isc_dsql_sql_info()

Returns requested information about a prepared DSQL statement.

```
Syntax  ISC_STATUS isc_dsql_sql_info(
        ISC_STATUS *status_vector,
        isc_stmt_handle *stmt_handle,
        unsigned short item_length,
        char *items,
        unsigned short buffer_length,
        char *buffer);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS</i> *	Pointer to the error status vector
<i>stmt_handle</i>	<i>isc_stmt_handle</i> *	Pointer to a statement handle previously allocated with <i>isc_dsql_allocate_statement()</i> or <i>isc_dsql_alloc_statement2()</i> ; the handle returns an error in <i>status_vector</i> if it is NULL
<i>item_length</i>	unsigned short	Number of bytes in the string of information items in <i>items</i>
<i>items</i>	<i>char</i> *	String of requested information items
<i>buffer_length</i>	unsigned short	Number of bytes in the result buffer, <i>buffer</i>
<i>buffer</i>	<i>char</i> *	User-provided buffer for holding returned data; must be large enough to hold the information requested

Description *isc_dsql_sql_info()* returns requested information about a statement prepared with a call to *isc_dsql_prepare()*. The main application need for this function is to determine the statement type of an unknown prepared statement, for example, a statement entered by the user at run time.

Requested information can include the:

- Statement type
- Number of input parameters required by the statement
- Number of output values returned by the statement
- Detailed information regarding each input parameter or output value, including its datatype, scale, and length
- The query plan prepared by the optimizer

Example The following illustrates a call to *isc_dsql_sql_info()* to determine the statement type of the statement whose handle is referenced by *stmt*:

```
int statement_type;
int length;
char type_item[] = {isc_info_sql_stmt_type};
char res_buffer[8];
isc_dsql_sql_info(
    status_vector,
    &stmt,
    /* Allocated previously by isc_dsql_allocate_statement() or
       isc_dsql_alloc_statement2() call. */
    sizeof(type_item),
    type_item,
    sizeof(res_buffer),
    res_buffer);

if (res_buffer[0] == isc_info_sql_stmt_type)
{
    length = isc_vax_integer(buffer[1], 2);
    statement_type = isc_vax_integer(buffer[3], length);
}
```

Return Value *isc_dsql_sql_info()* returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also ***isc_dsql_describe_bind()***, ***isc_dsql_describe()***, ***isc_vax_integer()***

For more information about determining unknown statement types at run time, see **“Determining an unknown statement type at runtime” on page 113 of Chapter 6, “Working with Dynamic SQL.”**

isc_encode_sql_date()

Translates a date from the C *struct tm* format to InterBase ISC_DATE format prior to inserting or updating a DATE value in a table.

Syntax

```
void isc_encode_sql_date(
    void *tm_date,
    ISC_DATE *ib_date);
```

Parameter	Type	Description
<i>tm_date</i>	<i>void *</i>	Pointer to a C <i>struct tm</i> structure
<i>ib_date</i>	<i>ISC_DATE *</i>	Pointer to a four-byte ISC_DATE structure containing a date in InterBase format

Description *isc_encode_sql_date()* translates a date in a C time structure into an ISC_DATE format internal to InterBase. This call is used prior to writing DATE data to a table to guarantee that the date is in a format recognized by InterBase.

Use the *isc_dsql* family of API calls to insert or update DATE data from the ISC_DATE structure in a table.

Note In InterBase 6, the DATE datatype is available only in dialect 3. It holds only date information, and does not include time information. In version 6 dialect 1, the TIMESTAMP datatype holds both date and time information and is exactly equivalent to the DATE datatype that was present in earlier versions of InterBase.

Example The following code fragment illustrates declaring time structures and calling *isc_encode_sql_date()* to translate a C time format into an InterBase date format prior to inserting or updating a table:

```
#include <time.h>
#include <ibase.h>
. . .
struct tm hire_time;
ISC_DATE hire_date;
. . .
/* Store date info into the tm struct here. */
. . .
isc_encode_sql_date(&hire_time, &hire_date);
/* Now use a DSQL INSERT or UPDATE statement to move the date into a
DATE column. */
```

Return Value None.

See Also [isc_decode_sql_date\(\)](#), [isc_encode_sql_time\(\)](#), [isc_encode_timestamp\(\)](#)

isc_encode_sql_time()

Translates a time from the C *struct tm* format to InterBase ISC_SQL_TIME format prior to inserting or updating a TIME value in a table.

Syntax

```
void isc_encode_sql_time(
    void *tm_date,
    ISC_TIME *ib_time);
```

Parameter	Type	Description
<i>tm_date</i>	<i>void *</i>	Pointer to a C <i>tm</i> structure
<i>ib_time</i>	<i>ISC_TIME *</i>	Pointer to a four-byte ISC_TIME structure containing a time in InterBase format

Description *isc_encode_sql_time()* translates a date in a C time structure into an ISC_TIME format internal to InterBase. This call is used prior to writing TIME data to a table to guarantee that the time is in a format recognized by InterBase.

Use the *isc_dsql* family of API calls to insert or update TIME data from the ISC_TIME structure in a table.

Example The following code fragment illustrates declaring time structures and calling *isc_encode_sql_time()* to translate a C time format into an InterBase date format prior to inserting or updating a table:

```
#include <time.h>
#include <ibase.h>
. . .
struct tm hire_time;
ISC_TIME hire_date;
. . .
/* Store time info into the tm struct here. */
. . .
isc_encode_sql_time(&hire_time, &hire_date);
/* Now use a DSQL INSERT or UPDATE statement to move the date into a
TIME column. */
```

Return Value None.

See Also [isc_decode_sql_time\(\)](#), [isc_encode_sql_date\(\)](#), [isc_encode_timestamp\(\)](#)

isc_encode_timestamp()

Translates a time from the C *struct tm* format to InterBase ISC_TIMESTAMP format prior to inserting or updating a TIMESTAMP value in a table.

Syntax

```
void isc_encode_timestamp(
    void *tm_date,
    ISC_TIMESTAMP *ib_timestamp);
```

Parameter	Type	Description
<i>tm_date</i>	<i>void *</i>	Pointer to a C <i>tm</i> structure
<i>ib_timestamp</i>	<i>ISC_TIMESTAMP *</i>	Pointer to an eight-byte ISC_TIMESTAMP structure containing a date and time in InterBase format

Description *isc_encode_timestamp()* translates a date in a C time structure into an ISC_TIMESTAMP format internal to InterBase. This call is used prior to writing TIMESTAMP data to a table to guarantee that the date and time are in a format recognized by InterBase. This call is exactly the same as the older *isc_encode_date()*, which is still available for backward compatibility.

Use the *isc_dsql* family of API calls to insert or update TIMESTAMP data from the ISC_TIMESTAMP structure in a table.

Example The following code fragment illustrates declaring time structures and calling *isc_encode_timestamp()* to translate a C time format into an InterBase date format prior to inserting or updating a table:

```
#include <time.h>
#include <ibase.h>
. . .
struct tm hire_time;
ISC_TIMESTAMP hire_date;
. . .
/* Store date and time info into the tm struct here. */
. . .
isc_encode_timestamp (&hire_time, &hire_date);
```

isc_event_block()

```
/* Now use a DSQL INSERT or UPDATE statement to move the date into a
TIMESTAMP column. */
```

Return Value None.

See Also [isc_decode_timestamp\(\)](#), [isc_encode_sql_date\(\)](#), [isc_encode_sql_time\(\)](#)

isc_event_block()

Allocates two event parameter buffers (EPBs) for subsequent use with other API event calls.

Syntax

```
long isc_event_block(
    char **event_buffer,
    char **result_buffer,
    unsigned short id_count,
    . . . );
```

Parameter	Type	Description
<i>event_buffer</i>	<i>char **</i>	Address of a character pointer; this function allocates and initializes an event parameter buffer and stores its address into the character pointer
<i>result_buffer</i>	<i>char **</i>	Address of a character pointer; this function allocates an event parameter buffer, and stores its address into the character pointer
<i>id_count</i>	<i>unsigned short</i>	Number of event identifier strings that follow
...	<i>char *</i>	Up to 15 null-terminated and comma-separated strings that each name an event

Description `isc_event_block()` must be called before any other event functions. It:

- Allocates two event parameter buffers of the same size, and stores their addresses into the character pointers addressed by *event_buffer* and *result_buffer*.
- Stores into the buffer referenced by *event_buffer* the names and event counts for each of the specified events. The names are the ones that appear as the final arguments to `isc_event_block()`. The event counts are initialized to zero and are used to specify how many times each event has been posted prior to each wait for events to occur.
- Returns the length, in bytes, of the buffers.

The buffers, and their lengths, are used in subsequent calls to the functions *isc_wait_for_event()*, *isc_que_events()*, and *isc_event_counts()*. *event_buffer* is used to indicate the events of interest, and to hold the counts in effect before a wait for one of the events. After an event is posted, *result_buffer* is filled in exactly as *event_buffer*, except that the event counts are updated. *isc_event_counts()* is then called to determine which events were posted between the time the counts were set in *event_buffer*, and the time the counts are set in *result_buffer*.

Example The following program fragment illustrates a call to *isc_event_block()*:

```
#define number_of_stocks 3;

char *event_buffer, *result_buffer;
long length;

length = isc_event_block(
    &event_buffer,
    &result_buffer,
    number_of_stocks,
    "DEC", "HP", "SUN");
```

Return Value *isc_event_block()* returns a number that is the size, in bytes, of each event parameter buffer it allocates.

See Also ***isc_event_counts()***, ***isc_que_events()***, ***isc_wait_for_event()***

isc_event_counts()

Compares event parameter buffers (EPBs) to determine which events have been posted, and prepares the event parameter buffers for the next call to *isc_que_events()* or *isc_wait_for_event()*.

Syntax

```
void isc_event_counts(
    ISC_STATUS *status_vector,
    short buffer_length,
    char *event_buffer,
    char *result_buffer);
```

Parameter	Type	Description
<i>status_vector</i>	<i>long *</i>	Pointer to the status vector, which is used to store the differences in event counts for each corresponding event in <i>event_buffer</i> and <i>result_buffer</i>
<i>buffer_length</i>	<i>short</i>	Length of the event parameter buffers, returned by the <i>isc_event_block()</i> call that allocated them
<i>event_buffer</i>	<i>char *</i>	Pointer to the event parameter buffer that specifies the event counts prior to the previous call to <i>isc_wait_for_event()</i> or <i>isc_que_events()</i>
<i>result_buffer</i>	<i>char *</i>	Pointer to the event parameter buffer filled in as a result of posting an event

Description *isc_event_counts()* compares the event counts in the event parameter buffers, *event_buffer* and *result_buffer*, and sets up to the first 15 elements of *status_array* to contain the differences. It then modifies *event_buffer* to contain the same event counts as *result_buffer* in preparation for the next call to either *isc_wait_for_event()* or *isc_que_events()*.

The counts in *event_buffer* specify how many times each event had been posted since the previous call to *isc_event_wait()* or *isc_que_events()*. The counts in *result_buffer* equal the values in *event_buffer* plus the number of additional times an event is posted after the current call to *isc_event_wait()* or *isc_que_events()*. If an event is posted after a call to either of these functions, its count is greater in *result_buffer* than in *event_buffer*. Other event counts may also be greater because an event may have been posted between calls to either of these functions. The values in *status_array* are the differences in values between *event_buffer* and *result_buffer*. This mechanism of comparing all the counts ensures that no event postings are missed.

Example The following program fragment illustrates the set-up and waiting on any of the events named “DEC”, “HP”, or “SUN”, then calling `isc_event_counts()` to determine which events have been posted:

```
#include <ibase.h>
#define number_of_stocks 3;

char *event_buffer, *result_buffer;
ISC_STATUS status_vector[20];
char *event_names[] = {"DEC", "HP", "SUN"};
long length;
int i;

length = isc_event_block(
    &event_buffer,
    &result_buffer,
    number_of_stocks,
    "DEC", "HP", "SUN");

isc_wait_for_event(
    status_vector,
    &database_handle, /* Set by previous isc_attach_database(). */
    length, /* Returned from isc_event_block(). */
    event_buffer,
    result_buffer);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector); /* Display error message. */
    return(1);
}

isc_event_counts(
    status_vector,
    (short) length,
    event_buffer,
    result_buffer);

for (i=0; i<number_of_stocks; i++)
    if (status_vector[i])
    {
```

isc_expand_dpb()

```
        /* The event has been posted. Do whatever is appropriate, for
example,
        initiating a buy or sell order. */
        ;
    }
```

Return Value None.

See Also [isc_que_events\(\)](#), [isc_wait_for_event\(\)](#)

isc_expand_dpb()

Dynamically builds or expands a database parameter buffer (DPB) to include database parameters.

Syntax

```
void isc_expand_dpb(
    char **dpb,
    unsigned short *dpb_size,
    . . . );
```

Parameter	Type	Description
<i>dpb</i>	<i>char **</i>	Pointer to an existing DPB
<i>dpb_size</i>	<i>unsigned short *</i>	Pointer to the current size, in bytes, of the DPB
...	<i>char *</i>	Pointers to items to insert into the expanded DPB

Description *isc_expand_dpb()* builds or expands a DPB dynamically. Its main use is to simplify the building of the DPB prior to a call to *isc_attach_database()*, or to allow an end user to supply a user name and password combination at run time. In many cases, the DPB must be constructed programmatically, but *isc_expand_dpb()* enables an application to pass user names, password, message file, and character set parameters to the function, which then adds them to an existing DPB.

A pointer to a previously allocated and initialized DPB must be passed to *isc_expand_dpb()* along with a pointer to a variable containing the current size of the DPB when this function is called. If the space allocated for the DPB is not large enough for the parameters passed to *isc_expand_dpb()*, then the function reallocates a larger DPB, preserving its current contents, and adds the new parameters.

To ensure proper memory management, applications that call *isc_expand_dpb()* should always allocate DPBs large enough to hold all anticipated parameters.

Example The following code calls `isc_expand_dpb()` to create a DPB, then attaches to a database using the newly created DPB. `user_name` and `user_password` are assumed to be variables whose values have been filled in, for example, after asking the user to specify the name and password to be used.

```
#include <ibase.h>
char *dpb;
ISC_STATUS status_vector[20];
isc_db_handle handle = NULL;
short dpb_length;

/* Build the database parameter buffer. */

dpb = (char *) malloc(50);
dpb_length = 0;

isc_expand_dpb(&dpb, &dpb_length, isc_dpb_user_name, user_name,
isc_dpb_password, user_password, NULL);

isc_attach_database(
    status_vector,
    0,
    "employee.db",
    &handle,
    dpb_length,
    dpb_buffer);
if (status_vector[0] == 1 && status_vector[1])
{
    /* An error occurred. */
    isc_print_status(status_vector);
    return(1);
}
```

Return Value None.

See Also **`isc_attach_database()`**

isc_get_segment()

Reads a segment from an open Blob.

Syntax `ISC_STATUS isc_get_segment(
 ISC_STATUS *status_vector,
 isc_blob_handle *blob_handle,
 unsigned short *actual_seg_length,
 unsigned short seg_buffer_length,
 char *seg_buffer);`

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>blob_handle</i>	<i>isc_blob_handle *</i>	Pointer to the handle of the Blob you want to read.
<i>actual_seg_length</i>	<i>unsigned short *</i>	Pointer to the actual segment length that InterBase reads into the buffer; useful if the segment length is shorter than the buffer length
<i>seg_buffer_length</i>	<i>unsigned short</i>	Length of the segment buffer
<i>seg_buffer</i>	<i>char *</i>	Pointer to the segment buffer

Description *isc_get_segment()* reads a Blob segment from a previously opened Blob. You can set the *seg_buffer_length* parameter to a size that is efficient for a particular type of Blob data. For example, if you are reading Blob data from a text file, you might set the segment buffer length to 80, to take advantage of the 72 to 80 character line lengths that are common in text files. By periodically checking the value of the actual segment length in your loop, you can determine an end-of-line or end-of-file condition.

Before reading any part of a Blob, you must open the Blob with a call to *isc_open_blob2()*. *isc_get_segment()* behaves differently depending on which call precedes it. If the most recent call is to *isc_open_blob2()*, then a call to *isc_get_segment()* reads the first segment in the Blob. If the most recent call is to *isc_get_segment()*, then it reads the next segment.

If Blob filters are specified when a Blob is opened, then each segment retrieved by *isc_get_segment()* is filtered on read.

Note Blob filters are not supported on NetWare.

You can read bitmaps and other binary files directly, without filtering, if you don't need to change from one format to another, say from .TIF to .JPEG. You can also store compressed bitmaps directly in a database in formats such as .JPG (JPEG), .BMP (Windows native bitmaps), or .GIF (CompuServe Graphic Interchange Format). No filtering is required.

You can store bitmaps in a database in row-major or column-major order.

If the buffer is not large enough to hold the entire current segment, the function returns *isc_segment*, and the next call to *isc_get_segment()* gets the next chunk of the oversized segment rather than getting the next segment.

When *isc_get_segment()* reads the last segment of the Blob, the function returns the code *isc_segstr_eof*.

For more information about reading data from a Blob, see **Chapter 7, “Working with Blob Data.”**

Example The following call gets a segment from one Blob and writes it to another:

```
get_status = isc_get_segment(status, &from_blob, &seg_len, 80,
buffer);
if (status[0] == 1 && status[1])
{
    isc_print_status(status);
    return(1);
}
if (get_status != isc_segstr_eof)
    write_status = isc_put_segment(status, &to_blob, seg_len, buffer);
if (status[0] == 1 && status[1])
{
    isc_print_status(status);
    return(1);
}
```

Return Value *isc_get_segment()* returns the second element of the status vector. Zero indicates success. *isc_segment* indicates the buffer is not large enough to hold the entire current segment; the next call to *isc_get_segment()* gets the next chunk of the oversized segment rather than getting the next segment. *isc_segstr_eof* indicates that the last segment of the Blob has been read. Any other nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also [isc_create_blob2\(\)](#), [isc_open_blob2\(\)](#), [isc_put_segment\(\)](#)

isc_interprete()

Extracts the text for an InterBase error message from the error status vector to a user-defined buffer.

Syntax

```
ISC_STATUS isc_interprete(
    char *buffer,
    ISC_STATUS **status_vector);
```

Parameter	Type	Description
<i>buffer</i>	<i>char*</i>	Application buffer for storing an InterBase error message
<i>status_vector</i>	<i>ISC_STATUS**</i>	Pointer to a pointer to the error status vector

Description Given both the location of a storage buffer allocated in a program, and the address of the status vector, *isc_interprete()* builds an error message string from the information in the status vector, puts the formatted string in the buffer where the program can manipulate it, and advances the status vector pointer to the start of the next cluster of error message information. For example, you might declare an error string buffer, call *isc_interprete()* to retrieve the first error message and insert the message into the buffer, write the buffer to a log file, then peek at the next cluster to see if it contains more error information.

isc_interprete() retrieves and formats a single message each time it is called. When an error occurs, however, the status vector usually contains more than one error message. To retrieve all relevant error messages, you must make repeated calls to *isc_interprete()* until no more messages are returned.

Note Do not pass the address of the status vector directly, because each time *isc_interprete()* is called, it modifies the pointer to the status vector to point to the start of the next available message.

To display all error messages on the screen instead of to a buffer, use *isc_print_status()*.

Example The following code declares a message buffer, a status vector, and a pointer to the vector, then illustrates how repeated calls are made to *isc_interprete()* to store all messages in the buffer:

```
#include <ibase.h>
char msg[512];
ISC_STATUS status_vector[20];
```

```

long *pvector; /* Pointer to pointer to status vector. */
FILE *efile; /* Code fragment assumes this points to an open file. */
. . .
pvector = status_vector; /* (Re)set to start of status vector. */
isc_interprete(msg, &pvector); /* Retrieve first message. */
fprintf(efile, "%s\n", msg); /* Write buffer to log file. */
msg[0] = '-'; /* Append leading hyphen to secondary messages. */
while(isc_interprete(msg + 1,&pvector)) /* More messages? */
{
    fprintf(efile, "%s\n", msg); /* If so, write them, too. */
}
fclose(efile);
. . .

```

Return Value If successful, `isc_interprete()` returns the length of the error message string it stores in *buffer*. It also advances the status vector pointer to the start of the next cluster of error message information.

If there are no more messages in the status vector, or if `isc_interprete()` cannot interpret the next message, it returns 0.

See Also [isc_print_sqlerror\(\)](#), [isc_print_status\(\)](#), [isc_sqlcode\(\)](#), [isc_sql_interprete\(\)](#)

isc_modify_user()

Modifies a user record from the password database, `isc4.gdb`.

Note Use of this function is deprecated. It is replaced by a full featured Services API. See [Chapter 12: “Working with Services” on page 199](#) and the reference entry for [“isc_service_start\(\)” on page 380](#).

Syntax

```

ISC_STATUS isc_modify_user(
    ISC_STATUS *status
    USER_SEC_DATA *user_sec_data);

```

Parameter	Type	Description
<i>status vector</i>	<code>ISC_STATUS *</code>	Pointer to the error status vector
<i>user_sec_data</i>	<code>USER_SEC_DATA *</code>	Pointer to a struct that is defined in <code>ibase.h</code>

Description The three security functions, `isc_add_user()`, `isc_delete_user()`, and `isc_modify_user()` mirror functionality that is available in the **gsec** command-line utility. `isc_modify_user()` modifies a record from **isc4.gdb**, InterBase's password database.

At a minimum, you must provide the user name. Any additional user information that you supply, such as first name, last name, or password, overwrites the information that is already in **isc4.gdb**.

If the server is not local, you must provide both a server name and a protocol. Valid choices for the protocol field are `sec_protocol_tcpip`, `sec_protocol_netbeui`, `sec_protocol_spx`, and `sec_protocol_local`.

InterBase reads the settings for the `ISC_USER` and `ISC_PASSWORD` environment variables if you do not provide a DBA user name and password.

The definition for the `USER_SEC_DATA` struct in **ibase.h** is as follows:

```
typedef struct {
    short  sec_flags;           /* which fields are specified */
    int    uid;                /* the user's id */
    int    gid;                /* the user's group id */
    int    protocol;           /* protocol to use for connection */
    char   *server;            /* server to administer */
    char   *user_name;         /* the user's name */
    char   *password;          /* the user's password */
    char   *group_name;        /* the group name */
    char   *first_name;        /* the user's first name */
    char   *middle_name;       /* the user's middle name */
    char   *last_name;         /* the user's last name */
    char   *dba_user_name;     /* the dba user name */
    char   *dba_password;      /* the dba password */
} USER_SEC_DATA;
```

When you pass this struct to one of the three security functions, you can tell it which fields you have specified by doing a bitwise OR of the following values, which are defined in **ibase.h**:

<code>sec_uid_spec</code>	0x01
<code>sec_gid_spec</code>	0x02
<code>sec_server_spec</code>	0x04
<code>sec_password_spec</code>	0x08
<code>sec_group_name_spec</code>	0x10
<code>sec_first_name_spec</code>	0x20
<code>sec_middle_name_spec</code>	0x40
<code>sec_last_name_spec</code>	0x80

```
sec_dba_user_name_spec      0x100
sec_dba_password_spec      0x200
```

No bit values are available for user name and password, since they are required.

The following error messages exist for this function:

Code	Value	Description
<i>isc_username_too_long</i>	335544747	The user name passed in is greater than 31 bytes
<i>isc_password_too_long</i>	335544748	The password passed in is longer than 8 bytes
<i>isc_username_required</i>	335544749	The operation requires a user name
<i>isc_password_required</i>	335544750	The operation requires a password
<i>isc_bad_protocol</i>	335544751	The protocol specified is invalid
<i>isc_dup_username_found</i>	335544752	The user name being added already exists in the security database.
<i>isc_username_not_found</i>	335544753	The user name was not found in the security database
<i>isc_error_adding_sec_record</i>	335544754	An unknown error occurred while adding a user
<i>isc_error_deleting_sec_record</i>	335544755	An unknown error occurred while deleting a user
<i>isc_error_modifying_sec_record</i>	335544756	An unknown error occurred while modifying a user
<i>isc_error_updating_sec_db</i>	335544757	An unknown error occurred while updating the security database

TABLE 13.20 Error messages for user security functions

Example The following example modifies **isc4.gdb** to change the password for the user Socks, using the bitwise OR technique for passing values from the USER_SEC_DATA struct.

```
{
  ISC_STATUS status[20];
  USER_SEC_DATA sec;

  sec.server      = "kennel";
  sec.dba_user_name= "sysdba";
  sec.dba_password = "masterkey";
  sec.protocol    = sec_protocol_tcpip;
  sec.user_name   = "socks";
```

isc_modify_user()

```
sec.password      = "feed_me!"; /* Note: do not hardcode passwords
*/
sec.sec_flags     = sec_server_spec
                  | sec_password_spec
                  | sec_dba_user_name_spec
                  | sec_dba_password_spec;

isc_add_user(status, &sec);
/* check status for errors */
if (status[0] == 1 && status[1])
{
    switch (status[1]) {
        case isc_username_too_long:
            printf("Security database cannot accept long user names\n");
            break;
        ...
    }
}
}
```

Return Value `isc_modify_user()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. See the “Description” section for this function for a list of error codes. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_add_user()`**, **`isc_delete_user()`**

isc_open_blob2()

Opens an existing Blob for retrieval and optional filtering.

```
Syntax  ISC_STATUS isc_open_blob2(
          ISC_STATUS *status_vector,
          isc_db_handle *db_handle,
          isc_tr_handle *trans_handle,
          isc_blob_handle *blob_handle,
          ISC_QUAD *blob_id,
          short bpb_length,
          char *bpb_address);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle *</i>	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>trans_handle</i>	<i>isc_tr_handle *</i>	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL
<i>blob_handle</i>	<i>isc_blob_handle *</i>	Pointer to the Blob handle, which must be NULL when you make this call
<i>blob_id</i>	<i>ISC_QUAD *</i>	Pointer to the 64-bit system-defined Blob ID, which is stored in a field in the table and points to the first segment of the Blob or to a page of pointers to Blob fragments
<i>bpb_length</i>	<i>short</i>	Length of the Blob parameter buffer (BPB)
<i>bpb_address</i>	<i>char *</i>	Pointer to the BPB

Description *isc_open_blob2()* opens an existing Blob for retrieval and optional filtering from one Blob subtype to another.

Note Using Blob filters is not supported on NetWare.

Input and output Blob filter types are passed to *isc_open_blob2()* as subtype information in a previously populated BPB, pointed to by *bpb_address*. If Blob filters are not needed or cannot be used, a BPB is not needed; pass 0 for *bpb_length* and NULL for *bpb_address*.

The *blob_id* identifies which particular Blob is to be opened. This *blob_id* is set by a sequence of DSQL function calls.

On success, *isc_open_blob2()* assigns a unique ID to *blob_handle*. Subsequent API calls use this handle to identify the Blob against which they operate.

After a blob is opened, its data can be read by a sequence of calls to *isc_get_segment()*.

When finished accessing the Blob, close it with *isc_close_blob()*.

For more information about opening a Blob for retrieval and optional filtering, see **Chapter 7, “Working with Blob Data.”**

Example The following fragment is excerpted from the example file, **api9.c**. The example program displays job descriptions that are passed through a filter.

```
while ((fetch_stat = isc_dsql_fetch(status, &stmt, 1, sqlda)) == 0)
{
    printf("\nJOB CODE: %5s GRADE: %d", job_code, job_grade);
    printf(" COUNTRY: %-20s\n\n", job_country);
    /* Open the blob with the fetched blob_id. */
    isc_open_blob2(status, &DB, &trans, &blob_handle, &blob_id, 9,
bpb);
    if (status[0] == 1 && status[1])
    {
        isc_print_status(status);
        return(1);
    }
}
```

Return Value *isc_open_blob2()* returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **isc_close_blob()**

isc_prepare_transaction()

Executes the first phase of a two-phase commit against multiple databases.

Syntax `ISC_STATUS isc_prepare_transaction(
 ISC_STATUS *status_vector,
 isc_tr_handle *trans_handle);`

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>trans_handle</i>	<i>isc_tr_handle *</i>	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL

Description *isc_prepare_transaction()* initiates the first phase of a two-phase commit under program direction. It alerts InterBase, which polls all database participants and waits for replies. The *isc_prepare_transaction()* function puts the transaction in limbo.

Because a call to this function indicates that you intend to control all phases of the commit, you must complete the second phase of the commit by explicitly calling the *isc_commit_transaction()* function.

If a call to *isc_prepare_transaction()* fails, the application should roll back the transaction with a call to the *isc_rollback_transaction()* function.

Note If you want InterBase to automatically perform the two-phase commit, call *isc_commit_transaction()* without calling *isc_prepare_transaction()*.

Example The following example executes the first phase of a two-phase commit and includes a rollback in case of failure:

```
isc_prepare_transaction(status_vector, &trans);
if (status_vector[0] == 1 && status_vector[1])
    rb_status = isc_rollback_transaction(status_vector, &trans)
else
{
    isc_commit_transaction(status_vector, &trans);
    if (!(status_vector[0] == 1 && status_vector[1]))
        fprintf(stderr, "Commit successful.\n");
}
```

isc_prepare_transaction2()

Return Value `isc_prepare_transaction2()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_commit_transaction()`, `isc_prepare_transaction2()`, `isc_rollback_transaction()`**

isc_prepare_transaction2()

Performs the first phase of a two-phase commit for multi-database transactions.

Syntax

```
ISC_STATUS isc_prepare_transaction2(  
    ISC_STATUS *status_vector,  
    isc_tr_handle *trans_handle,  
    unsigned short msg_length,  
    char *message);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>trans_handle</i>	<i>isc_tr_handle *</i>	Pointer to a transaction handle whose value has been set by a previous <code>isc_start_transaction()</code> call; <i>trans_handle</i> returns an error if NULL
<i>msg_length</i>	<i>unsigned short</i>	Length of message in bytes
<i>message</i>	<i>char *</i>	Transaction description buffer

Description `isc_prepare_transaction2()` performs the first phase of a two-phase commit, just as `isc_prepare_transaction()` does, but `isc_prepare_transaction2()` expects you to provide two additional arguments:

- An information message to write to the RDB\$TRANSACTION_DESCRIPTION column in the RDB\$TRANSACTIONS system table that describes the transaction to commit, so that recovery is possible in the event a system crash occurs during the completion of the commit.
- The length, in bytes, of the information message.

By electing to use `isc_prepare_transaction2()`, you are, in effect, disabling the automatic recovery functions inherent in the two-phase commit. It is your responsibility to deal with recovery issues that might occur during failure of the two-phase commit. Normally, InterBase automatically writes to the `RDB$TRANSACTION_DESCRIPTION` column in the `RDB$TRANSACTIONS` system table information that makes it possible to reconnect following a system crash during the commit. You can manually write a message string into `RDB$TRANSACTIONS`, by using the *message* parameter in this function.

At the risk of preventing recovery in the event of a system crash, you might choose to avoid writing a message to `RDB$TRANSACTION` altogether if you determine that there is too much overhead associated with this extra action every time your application commits.

Example The following example executes the first phase of a two-phase commit and includes a rollback in case of failure:

```
isc_prepare_transaction2(status_vector, &trans, msg_len, msg);
if (status_vector[0] == 1 && status_vector[1])
    rb_status = isc_rollback_transaction(status_vector, &trans);
```

Return Value `isc_prepare_transaction2()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_commit_transaction()`, `isc_prepare_transaction()`, `isc_rollback_transaction()`**

isc_print_sqlerror()

Displays an *SQLCODE* value, a corresponding SQL error message, and any additional InterBase error messages in the error status vector.

Syntax void isc_print_sqlerror(
 short SQLCODE,
 ISC_STATUS *status_vector);

Parameter	Type	Description
<i>SQLCODE</i>	<i>short</i>	Variable containing an SQLCODE value
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector

Description During the processing of DSQL API calls, SQL errors can occur. SQL errors are generally reported in a variable called *SQLCODE*. DSQL calls return error information to a user-defined error status vector like any other API call, but *isc_print_sqlerror()* can be used to interpret the primary error condition as an SQL error message for direct display on the screen. To use *isc_print_sqlerror()*, an application must declare both an *SQLCODE* variable for holding the SQL error number, and an error status vector for holding InterBase error information. *isc_print_sqlerror()* displays the *SQLCODE* value, a related SQL error message, and any additional InterBase error messages in the status array.

Note Some windowing systems do not permit direct screen writes. Do not use *isc_print_sqlerror()* when developing applications for these environments. Instead, use *isc_sql_interprete()* and *isc_interprete()* to capture messages to a buffer for display.

Example The following code calls *isc_print_sqlerror()* when an error occurs:

```
#include <ibase.h>
long SQLCODE;
ISC_STATUS status_vector[20];
. . .
if (status_vector[0] == 1 && status_vector[1])
{
    SQLCODE = isc_sqlcode(status_vector);
    isc_print_sqlerror(SQLCODE, status_vector);
}
```

Return Value None.

See Also [isc_interprete\(\)](#), [isc_print_status\(\)](#), [isc_sql_interprete\(\)](#), [isc_sqlcode\(\)](#)

isc_print_status()

Builds and displays error messages based on the contents of the InterBase error status vector.

Syntax `ISC_STATUS isc_print_status(ISC_STATUS *status_vector);`

Parameter	Type	Description
<i>status_vector</i>	<code>ISC_STATUS *</code>	Pointer to the error status vector

Description `isc_print_status()` builds all error messages based on the contents of the error status vector, and displays them on the screen. *status_vector* must be declared in the program as an array of twenty elements.

Example The following code displays error messages when an error occurs during processing:

```
#include <ibase.h>
ISC_STATUS status_vector[20];
. . .
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector);
    return(1);
}
```

Return Value `isc_print_status()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_interprete()`, `isc_print_sqlerror()`, `isc_sqlcode()`, `isc_sql_interprete()`**

isc_put_segment()

Writes a Blob segment.

Syntax

```
ISC_STATUS isc_put_segment(
    ISC_STATUS *status_vector,
    isc_blob_handle *blob_handle,
    unsigned short seg_buffer_length,
    char *seg_buffer);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>blob_handle</i>	<i>isc_blob_handle *</i>	Pointer to the handle of the Blob to which you want to write; use <i>isc_create_blob2()</i> to set a value for this handle
<i>seg_buffer_length</i>	<i>unsigned short</i>	Length of the Blob segment buffer
<i>seg_buffer_address</i>	<i>char *</i>	Pointer to the Blob segment buffer that contains data for writing

Description *isc_put_segment()* writes a Blob segment in *seg_buffer_address* to a Blob previously created and opened with *isc_create_blob2()*.

If a Blob filter was specified when the Blob was created, then each segment is filtered before storing the result into the Blob.

The behavior of *isc_put_segment()* depends on what call preceded it. If the most recent call was to *isc_create_blob()* or *isc_create_blob2()*, then a call to *isc_put_segment()* writes the first segment of the Blob. If the most recent call was to *isc_put_segment()*, then it writes the next segment.

You can write bitmaps and other binary files directly, without filtering, unless you intend to change from one format to another, say from .GEM to .BMP. You can also store compressed bitmaps directly in a database, in formats such as .JPG (JPEG), .BMP (Windows native bitmaps), or .GIF (CompuServe Graphic Interchange Format).

You can store bitmaps in your database in row-major or column-major order.

You cannot update a Blob directly. If you want to modify Blob data, you must do one of the following:

- Create a new Blob.
- Read the old Blob data into a buffer where you can edit or modify it.

- Write the modified data to the new Blob.
- Prepare and execute an UPDATE statement that will modify the Blob column to contain the Blob ID of the new Blob, replacing the old Blob's Blob ID.

For more information about creating and writing Blob data, see **Chapter 7, “Working with Blob Data.”**

Note To read a segment that you wrote with a call to `isc_put_segment()`, you must close the Blob with `isc_close_blob()`, and then open it with `isc_open_blob2()`.

Example The following example reads a segment of one Blob and writes it to another Blob:

```
get_status = isc_get_segment(status, &from_blob, &seg_len, 80,
buffer);
if (status[0] == 1 && status[1])
{
    isc_print_status(status);
    return(1);
}
if (get_status != isc_segstr_eof)
    write_status = isc_put_segment(status, &to_blob, seg_len, buffer);
if (status[0] == 1 && status[1])
{
    isc_print_status(status);
    return(1);
}
```

Return Value `isc_put_segment()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_close_blob()`, `isc_get_segment()`, `isc_open_blob2()`**

isc_que_events()

Requests asynchronous notification of one of a specified group of events.

Syntax

```
ISC_STATUS isc_que_events(
    ISC_STATUS *status_vector,
    isc_db_handle *db_handle,
    ISC_LONG *event_id,
    short length,
    char *event_buffer,
    isc_callback event_function,
    void *event_function_arg);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle *</i>	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> ; the handle identifies the database against which the events are expected to be posted <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>event_id</i>	<i>ISC_LONG *</i>	Pointer to an event identifier to set
<i>length</i>	<i>short</i>	Length of the event parameter buffers, returned by the <i>isc_event_block()</i> call which allocated them
<i>event_buffer</i>	<i>char *</i>	Pointer to the event parameter buffer that specifies the current counts of the events to be waited on; this buffer should have been initially allocated and filled in by a call to <i>isc_event_block()</i>
<i>event_function</i>	<i>isc_callback</i>	Pointer to the address of the function to receive event notification
<i>event_function_arg</i>	<i>void *</i>	First argument to be passed to <i>event_function</i> , usually a pointer to the event parameter buffer you want filled in with updated event counts

Description `isc_que_events()` is called to request asynchronous notification of any of the events listed in `event_buffer`. Upon completion of the call, but before events are posted, control is returned to the calling application, which can continue other processing. When a requested event is posted, InterBase calls the function specified in `event_function` to process event occurrence.

After `event_function` is called, you must call `isc_que_events()` again if you want to start another asynchronous wait on the specified events.

Note `isc_que_events()` cannot be called from within `event_function`.

If you want to cancel your `isc_que_events()` request for asynchronous event notification, call `isc_cancel_events()`.

Note To request *synchronous* notification, call `isc_wait_for_event()`.

Example The following program fragment illustrates calling `isc_que_events()` to wait asynchronously for event occurrences. Within a loop, it performs other processing, and checks the event flag (presumably set by the specified event function) to determine when an event has been posted. If one has, the program resets the event flag, calls `isc_event_counts()` to determine which events have been posted since the last call to `isc_que_events()`, and calls `isc_que_events()` to initiate another asynchronous wait.

```
#include <ibase.h>
#define number_of_stocks 3;
#define MAX_LOOP 10

char *event_names[] = {"DEC", "HP", "SUN"};
char *event_buffer, *result_buffer;
ISC_STATUS count_array[number_of_stocks];
short length;
ISC_LONG event_id;
int i, counter;
int event_flag = 0;

length = (short)isc_event_block(
    &event_buffer,
    &result_buffer,
    number_of_stocks,
    "DEC", "HP", "SUN");

isc_que_events(
    status_vector,
    &database_handle, /* Set in previous isc_attach_database(). */
    &event_id,
```

```
isc_que_events()
```

```
    length, /* Returned from isc_event_block(). */
    event_buffer,
    (isc_callback)event_function,
    result_buffer);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector); /* Display error message. */
    return(1);
};

counter = 0;
while (counter < MAX_LOOP)
{
    counter++;

    if (!event_flag)
    {
        /* Do whatever other processing you want. */
        ;
    }
    else
    {
        event_flag = 0;
        isc_event_counts(
            count_array,
            length,
            event_buffer,
            result_buffer);
        if (status_vector[0] == 1 && status_vector[1])
        {
            isc_print_status(status_vector); /* Display error message.
*/
            return(1);
        }

        for (i=0; i<number_of_stocks; i++)
            if (count_array[i])
            {
                /* The event has been posted. Do whatever is appropriate,
                for example, initiating a buy or sell order.
                Note: event_names[i] tells the name of the event
                corresponding to count_array[i]. */
                ;
            }
    }
}
```

```

    }

    isc_que_events(
        status_vector,
        &database_handle,
        &event_id,
        length,
        event_buffer,
        (isc_callback)event_function,
        result_buffer);
    if (status_vector[0] == 1 && status_vector[1])
    {
        isc_print_status(status_vector); /* Display error message.
*/
        return(1);
    }
} /* End of else. */
} /* End of while. */

/* Let InterBase know you no longer want to wait asynchronously. */
isc_cancel_events(
    status_vector,
    &database_handle,
    &event_id);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector); /* Display error message. */
    return(1);
}

```

Return Value `isc_que_events()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_cancel_events()`, `isc_event_block()`, `isc_event_counts()`, `isc_wait_for_event()`**

For more information about writing an asynchronous event trap (AST) function, see **Chapter 11, “Working with Events.”**

isc_rollback_retaining()

Undoes changes made by a transaction and retains the transaction context after the rollback.

Syntax `ISC_STATUS isc_rollback_retaining(
 ISC_STATUS *status_vector,
 isc_tr_handle *trans_handle);`

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>trans_handle</i>	<i>isc_tr_handle *</i>	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; this function returns an error if <i>trans_handle</i> is NULL

Description *isc_rollback_retaining()* rolls back an active transaction and immediately clones itself. This means that the function retains the transaction name, system resources associated with the transaction, and the current state of any open cursors in the transaction. Although the function is actually initiating a new transaction, by assigning the new transaction the existing transaction handle it is, in effect, keeping the transaction open after the rollback. This results in improved performance by allowing an application to minimize the overhead of initiating additional transactions. *isc_rollback_retaining()* allows you to roll back updates while keeping a cursor open.

You can initiate a rollback within the active transaction but the rollback only affects uncommitted updates. In other words, a rollback is legal, even after the transaction context has been passed to the cloned transaction, but, in that case, the rollback will only affect the updates your application has made to the database since the last commit or rollback.

To audit the rollbacks made by your calls to this function, check the first element in the status vector to see if the call was successful. If this element contains a zero, the call was successful.

The transaction ends when you commit or roll back without using the retention feature, with a call to *isc_commit_transaction()* or *isc_rollback_transaction()*.

Examples The following C/C++ code rolls back a transaction, prints a message, and starts a new transaction with the same handle within the same request:

```
if (!isc_rollback_retaining(status, &retained_trans))
{
    fprintf(stderr, "Rolled back and retained\n");
    isc_print_status(status);
}
```

The following C/C++ code rolls back a transaction, prints a confirmation message, starts a new transaction with the same handle within the same request, or, if the rollback fails, prints an error message and rolls back.

```
isc_rollback_retaining(status, &retained_trans);
if (status[0] == 1 && status[1])
{
    fprintf(stderr, "Error retaining; rolling back instead.\n");
    rb_status = isc_rollback_transaction(status, &retained_trans);
}
else
{
    fprintf(stderr, "Rollback retaining successful.\n");
    tr_count++; /* Increments the number of recycles. */
}
```

Return Value `isc_rollback_retaining()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_commit_retaining()`, `isc_commit_transaction()`, `isc_rollback_transaction()`, `isc_start_transaction()`**

isc_rollback_transaction()

Undoes changes made by a transaction, and restores the database to its state prior to the start of the specified transaction.

Syntax `ISC_STATUS isc_rollback_transaction(
ISC_STATUS *status_vector,
isc_tr_handle *trans_handle);`

Parameter	Type	Description
<i>status_vector</i>	<code>ISC_STATUS *</code>	Pointer to the error status vector
<i>trans_handle</i>	<code>isc_tr_handle *</code>	Pointer to a transaction handle whose value has been set by a previous <code>isc_start_transaction()</code> call; <i>trans_handle</i> returns an error if NULL

Description `isc_rollback_transaction()` rolls back a specified transaction, closes record streams, frees system resources, and sets the transaction handle to zero. It is typically used to undo all database changes made by a transaction when an error occurs.

A call to this function can fail only if:

- You pass a NULL or invalid transaction handle.
- The transaction dealt with more than one database and a communications link fails during the rollback operation. If that happens, subtransactions on the remote node will end up in limbo. You must use the database maintenance utility to manually roll back those transactions.

Example The following call rolls back a transaction:

```
isc_rollback_transaction(status_vector, &trans);
```

Return Value `isc_rollback_transaction()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_commit_transaction()`, `isc_rollback_retaining()`, `isc_start_transaction()`**

isc_service_attach()

Attaches to the InterBase Services Manager facility. You must do this before using the InterBase services functions to request execution of tasks or query information from the Services Manager.

Syntax

```
ISC_STATUS isc_service_attach(
    ISC_STATUS *status_vector,
    unsigned short service_length,
    char *service,
    isc_svc_handle *svc_handle,
    unsigned short spb_length,
    char *spb);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>service_length</i>	<i>unsigned short</i>	Length in characters of the service name; a value of zero means that the service name is a null-terminated string
<i>service</i>	<i>char *</i>	String containing the name of the service to which the client requests an attachment
<i>svc_handle</i>	<i>isc_svc_handle *</i>	Pointer to a long value containing the handle of the service structure
<i>spb_length</i>	<i>unsigned short</i>	Length in bytes of the services parameter buffer
<i>spb</i>	<i>char *</i>	Pointer to a services parameter buffer

Description You can use this function to attach to the Services Manager on a given InterBase server. The InterBase service must be running on that host before you can attach to the Services Manager.

You must specify the hostname and the literal string *service_mgr* in the *service* argument. For example, *jupiter:service_mgr* is the string you use to connect to the Services Manager on host *jupiter* using TCP/IP as the network protocol.

You must specify a user ID and the corresponding password as part of the options in the service parameter buffer. The Services Manager uses this user ID when performing service tasks you request.

There are components in the InterBase Express™ package for Delphi and C++Builder that provide a visual interface to the Services Manager. See the *Developer's Guide*.

isc_service_detach()

Example See “**Attaching to the Services Manager with isc_service_attach()**” on page 202 for an example using C/C++ code.

Return Value `isc_service_attach()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_service_detach()`, `isc_service_query()`, `isc_service_start()`**

isc_service_detach()

Terminates the attachment to the InterBase Services Manager.

Syntax

```
ISC_STATUS isc_service_detach(  
    ISC_STATUS *status_vector,  
    isc_svc_handle *svc_handle);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS</i> *	Pointer to the error status vector
<i>svc_handle</i>	<i>isc_svc_handle</i> *	Pointer to a long value containing the handle of the service structure

Description After you have performed all tasks and retrieved all information needed from the Services Manager, you should use this function to detach.

There are components in the InterBase Express™ package for Delphi and C++Builder that provide a visual interface to the Services Manager. See the *Developer's Guide*.

Example See “**Detaching from a Services Manager with isc_service_detach()**” on page 203 for an example using C/C++ code.

Return Value `isc_service_detach()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_service_attach()`**, **`isc_service_query()`**, **`isc_service_start()`**

isc_service_query()

Requests and retrieves information about the InterBase server to which the client is attached.

Syntax

```
ISC_STATUS isc_service_query(
    ISC_STATUS *status_vector,
    isc_svc_handle *svc_handle,
    isc_resv_handle *reserved,
    unsigned short send_spb_length,
    char *send_spb,
    unsigned short request_spb_length,
    char *request_spb,
    unsigned short buffer_length,
    char *buffer);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>svc_handle</i>	<i>isc_svc_handle *</i>	Pointer to a long value containing the handle of the service structure
<i>reserved</i>	<i>isc_resv_handle *</i>	Reserved for future use; should be NULL
<i>send_spb_length</i>	<i>unsigned short</i>	Length in bytes of the service parameter buffer
<i>send_spb</i>	<i>char *</i>	Pointer to a service parameter buffer containing flags for the Services Manager
<i>request_spb_length</i>	<i>unsigned short</i>	Length in bytes of the request buffer

isc_service_query()

Parameter	Type	Description
<i>request_spb</i>	<i>char *</i>	Pointer to a buffer containing item specifiers for requested information
<i>buffer_length</i>	<i>unsigned short</i>	Length in bytes of the return buffer
<i>buffer</i>	<i>char *</i>	Pointer to a buffer containing information received from the Services Manager

Description Use *isc_service_query()* to request information from the Services Manager. You must have an active connection to a running Services Manager, made using *isc_service_attach()* (see **page 376**).

There are components in the InterBase Express™ package for Delphi and C++Builder that provide a visual interface to the Services Manager. See the *Developer's Guide*.

Example There are several examples of using *isc_service_query()* with C/C++ in **“Querying the Services Manager” on page 220**.

Return Value *isc_service_query()* returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also ***isc_service_attach()***, ***isc_service_detach()***, ***isc_service_start()***

isc_service_start()

Performs a service task on the InterBase server to which the client is attached.

Syntax `ISC_STATUS isc_service_start(
 ISC_STATUS *status_vector,
 isc_svc_handle *svc_handle,
 isc_resv_handle *reserved,
 unsigned short spb_length,
 char *spb);`

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>svc_handle</i>	<i>isc_svc_handle *</i>	Pointer to a long value containing the handle of the service structure
<i>reserved</i>	<i>isc_resv_handle *</i>	Reserved for future use; should be NULL
<i>spb_length</i>	<i>unsigned short</i>	Length in bytes of the service parameter buffer
<i>spb</i>	<i>char *</i>	Pointer to a service parameter buffer containing flags and optional arguments instructing the Services Manager to perform specified tasks

Description Use *isc_service_start()* to initiate a task execution by the Services Manager. You must have an active connection to a running Services Manager, made using *isc_service_attach()* (see [page 376](#)).

There are components in the InterBase Express™ package for Delphi and C++Builder that provide a visual interface to the Services Manager. See the *Developer's Guide*.

Example There are several examples of using *isc_service_start()* with C/C++ in **“Invoking service tasks with *isc_service_start()*”** on [page 204](#).

Return Value *isc_service_start()* returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also ***isc_service_attach()***, ***isc_service_detach()***, ***isc_service_query()***

isc_sqlcode()

Translates an InterBase error code in the error status vector to an SQL error code number.

Syntax `ISC_LONG isc_sqlcode (ISC_STATUS *status_vector);`

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS</i> *	Pointer to the error status vector

Description *isc_sqlcode()* searches *status_vector* for a reported SQL error, and if it finds it, translates the InterBase error code number into an appropriate SQL error code. Typically, this call is used to populate a program variable (usually called *SQLCODE* for portability among SQL implementations) with an SQL error number for use in an SQL error-handling routine.

Example The following code illustrates how *isc_sqlcode()* might be called in a DSQL application:

```
#include <ibase.h>
long SQLCODE;
ISC_STATUS status_vector[20];
. . .
if (status_vector[0] == 1 && status_vector[1])
{
    SQLCODE = isc_sqlcode(status_vector);
    isc_print_sqlerror(SQLCODE, status_vector);
}
```

Return Value If successful, *isc_sqlcode()* returns the first valid SQL error code decoded from the InterBase status vector.

If no valid SQL error code is found, *isc_sqlcode()* returns -999.

See Also [isc_interprete\(\)](#), [isc_print_sqlerror\(\)](#), [isc_print_status\(\)](#), [isc_sql_interprete\(\)](#)

isc_sql_interprete()

Builds an SQL error message string and stores it in a user-defined buffer.

Syntax void *isc_sql_interprete*(
 short *SQLCODE*,
 char **buffer*,
 short *buffer_length*);

Parameter	Type	Description
<i>SQLCODE</i>	<i>short</i>	Variable containing an <i>SQLCODE</i> value
<i>buffer</i>	<i>char*</i>	Application buffer into which to store an SQL error message
<i>buffer_length</i>	<i>short</i>	Length, in bytes, of <i>buffer</i>

Description Given an *SQLCODE* value less than zero, *isc_sql_interprete()* builds a corresponding SQL error message string, and stores it in a user-defined buffer. The size of the buffer, in bytes, must also be passed to this function.

To display an SQL error message corresponding to an *SQLCODE* value, use *isc_print_sqlerror()* instead of this call.

Example The following code fragment illustrates a call to *isc_sql_interprete()*:

```
#include <ibase.h>
long SQLCODE;
char err_buf[256];
. . .
if (status_vector[0] == 1 && status_vector[1])
{
    SQLCODE = isc_sqlcode(status_vector);
    isc_sql_interprete(SQLCODE, err_buf, sizeof(err_buff));
}
```

Return Value None.

See Also [isc_interprete\(\)](#), [isc_print_sqlerror\(\)](#), [isc_print_status\(\)](#), [isc_sqlcode\(\)](#)

isc_start_multiple()

Begins a new transaction against multiple databases.

Syntax `ISC_STATUS isc_start_multiple(
 ISC_STATUS *status_vector,
 isc_tr_handle *trans_handle,
 short db_handle_count,
 void *teb_vector_address);`

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>trans_handle</i>	<i>isc_tr_handle *</i>	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL
<i>db_handle_count</i>	<i>short</i>	Number of database handles passed in this call via <i>transaction existence buffers</i> (TEBs)
<i>teb_vector_address</i>	<i>void *</i>	Pointer to the TEB

Description Call *isc_start_multiple()* if you:

- Are using a language that does not support a variable number of arguments in a function call.
- Do not know how many databases you want to attach to when coding the start transaction function.

isc_start_multiple() passes information about each target database to InterBase. That information is stored in an array of transaction existence blocks (TEBs) pointed to by the *teb_vector* parameter.

teb_vector is a pointer to a byte array that consists of consecutive TEBs, one TEB for each database to connect to. Each TEB consists of three items: a pointer to the database handle for a database against which the transaction should run; the length, in bytes, of the transaction parameter buffer (TPB) for the database, and a pointer to the TPB. The items in a TEB correspond to the items passed directly as parameters in calls to *isc_start_transaction()*. C programmers should use *isc_start_transaction()* instead of *isc_start_multiple()* whenever possible because it does not require setting up TEBs.

For more information about establishing TEBs and calling *isc_start_multiple()*, see **“Calling *isc_start_multiple()*” on page 71 of Chapter 5, “Working with Transactions.”**

Example The following program starts a multiple-database transaction:

```
#include <ibase.h>

typedef struct { /* Define the ISC_TEB structure. */
    int *dbb_ptr;
    long tpb_len;
    char *tpb_ptr;
} ISC_TEB;

ISC_TEB teb_vec[2]; /* Declare the TEB vector. */

ISC_STATUS isc_status[20]; /* Status vector. */
long *db0, *db1, /* Database handle. */
long *trans; /* Transaction handle. */

static char
    isc_tpb_0[] = { /* Declare the first transaction parameter
buffer. */
        isc_tpb_version3, /* InterBase version. */
        isc_tpb_write, /* Read-write access. */
        isc_tpb_consistency, /* Serializable. */
        isc_tpb_wait, /* Wait on lock. */
        isc_tpb_lock_write, 3, /* Reserving IDS for update. */
        'I', 'D', 'S',
        isc_tpb_protected}, /* Don't allow other transactions to
write to the table. */

    isc_tpb_1[] = { /* Declare the second transaction.*/
        /* Parameter buffer. */
        isc_tpb_version3, /* InterBase version. */
        isc_tpb_write, /* Read-write access. */
        isc_tpb_consistency, /* Serializable. */
        isc_tpb_wait, /* Wait on lock. */
        isc_tpb_lock_write, 3, /* Reserving table OZS for update. */
        'O', 'Z', 'S',
        isc_tpb_protected}; /* Don't allow other transactions to
write to the table. */

main()
{
    db0 = db1 = 0;
```

isc_start_multiple()

```
trans = 0;

/* If you can't attach to test_0 database, attach to test_1. */

isc_attach_database(isc_status, 0, "test_0.gdb", &db0, 0,0);
if (isc_status[0] == 1 && isc_status[1])
    isc_attach_database(isc_status, 0, "test_1.gdb", &db1, 0,0);

if (db0 && db1)
{
    /* Assign database handles, tpb length, and
       tpb handle to the teb vectors. */
    teb_vec[0].dbb_ptr = &db0;
    teb_vec[0].tpb_len = sizeof (isc_tpb_0);
    teb_vec[0].tpb_ptr = isc_tpb_0;

    teb_vec[1].dbb_ptr = &db1;
    teb_vec[1].tpb_len = sizeof (isc_tpb_1);
    teb_vec[1].tpb_ptr = isc_tpb_1;

    if (isc_start_multiple(isc_status, &trans, 2, teb_vec))
        isc_print_status(isc_status);
}

if (trans)
    isc_commit_transaction(isc_status, &trans);

if (db0 && !trans)
    isc_detach_database(isc_status, &db0);

if (db1 && !(trans && db0))
    isc_detach_database(isc_status, &db1);

if (isc_status[0] == 1 && isc_status[1])
    isc_print_status(isc_status);
}
```

Return Value `isc_start_multiple()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **isc_commit_transaction(), isc_prepare_transaction(), isc_prepare_transaction2(), isc_rollback_transaction(), isc_start_transaction()**

For more information about transaction handles, see **“Creating transaction handles” on page 61 of Chapter 5, “Working with Transactions.”** For more information about creating and populating a TPB, see **“Creating a transaction parameter buffer” on page 62 of Chapter 5, “Working with Transactions.”** For more information on TEBs, see **“Calling isc_start_multiple()” on page 71 of Chapter 5, “Working with Transactions.”**

isc_start_transaction()

Starts a new transaction against one or more databases.

Syntax

```
ISC_STATUS isc_start_transaction(
    ISC_STATUS *status_vector,
    isc_tr_handle *trans_handle,
    short db_handle_count,
    isc_db_handle *db_handle,
    unsigned short tpb_length,
    char *tpb_address,
    [isc_db_handle *db_handle,
    unsigned short tpb_length,
    char *tpb_address ...]);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>trans_handle</i>	<i>isc_tr_handle *</i>	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL
<i>db_handle_count</i>	<i>short</i>	Number of database handles passed in this call

isc_start_transaction()

Parameter	Type	Description
<i>db_handle</i>	<i>isc_db_handle *</i>	Pointer to a database handle set by a previous call to <i>isc_attach_database()</i> ; the handle identifies the database against which the events are expected to be posted <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>tpb_length</i>	<i>unsigned short</i>	Length of the <i>transaction parameter buffer</i> (TPB)
<i>tpb_address</i>	<i>char *</i>	Pointer to the TPB

Description *isc_start_transaction()* starts a new transaction against one or more databases specified as database handles.

Note If you have a variable number of databases to update, or are using a language that does not support a variable number of arguments in a function call, use *isc_start_multiple()* instead of *isc_start_transaction()*.

A single transaction can access multiple databases. This function passes information about each database it accesses and the conditions of access for that database in a *transaction parameter buffer* (TPB). The TPB is a variably-sized vector of bytes declared and populated by the program. It contains information describing intended transaction behavior such as its access and lock modes.

isc_start_transaction() can start a transaction against up to 16 databases. You must pass a database handle and a TPB for *each* referenced database. If you want to use defaults for the transaction, set *tpb_length* to zero. In this case, *tpb_vector* is a NULL pointer.

Example The following program includes a call to the start transaction function:

```
#include <ibase.h>

long
isc_status[20], /* Status vector. */
*db, /* Database handle. */
*trans, /* Transaction handle. */

static char
isc_tpb_0[] = {
    isc_tpb_version3, /* InterBase version. */
    isc_tpb_write, /* Read-write access. */
    isc_tpb_consistency, /* Consistency-mode transaction. */
    isc_tpb_wait, /* Wait on lock. */
    isc_tpb_lock_write, 3, /* Reserving IDS table for update. */
```

```

        "I", "D", "S",
        isc_tpb_protected}; /* Don't allow other transactions to
                               write against this table. */
main()
{
    db = trans = 0;
    isc_attach_database(isc_status, 0, "test.gdb", &db, 0,0);

    if (db)
    {
        isc_start_transaction(
            isc_status, &trans, 1, &db,
            sizeof(isc_tpb_0), isc_tpb_0);
        if (isc_status[0] == 1 && isc_status[1])
            isc_print_status(isc_status);
    }
    if (trans)
        isc_commit_transaction(isc_status, &trans);

    if (db && !trans)
        isc_detach_database(isc_status, &db);

    if (status_vector[0] == 1 && status_vector[1])
        isc_print_status(isc_status);
}

```

Return Value `isc_start_transaction()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_commit_transaction()`, `isc_prepare_transaction()`, `isc_prepare_transaction2()`, `isc_rollback_transaction()`, `isc_start_multiple()`**

For more information about transaction handles, see **“Creating transaction handles” on page 61 of Chapter 5, “Working with Transactions.”** For more information about creating and populating a TPB, see **“Creating a transaction parameter buffer” on page 62 of Chapter 5, “Working with Transactions.”**

isc_transaction_info()

Returns information about the specified named transaction.

Syntax

```
ISC_STATUS isc_transaction_info(
    ISC_STATUS *status_vector,
    isc_tr_handle *trans_handle,
    short item_list_buffer_length,
    char *item_list_buffer,
    short result_buffer_length,
    char *result_buffer);
```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>trans_handle</i>	<i>isc_tr_handle *</i>	Pointer to a transaction handle whose value has been set by a previous <i>isc_start_transaction()</i> call; <i>trans_handle</i> returns an error if NULL
<i>item_list_buffer_length</i>	<i>short</i>	Number of bytes in the item-list buffer
<i>item_list_buffer</i>	<i>char *</i>	Pointer to the item-list buffer
<i>result_buffer_length</i>	<i>short</i>	Number of bytes in the result buffer
<i>result_buffer</i>	<i>char *</i>	Pointer to the result buffer

Description *isc_transaction_info()* returns information necessary for keeping track of transaction IDs. This call is used internally by *isc_prepare_transaction()*. You should not need to use it in your own applications.

You can explicitly retrieve information about the transaction ID by including the following constant in the item-list buffer, where the transaction items about which you want information are listed:

Item	Purpose	Size of next value	Value
<i>isc_info_tra_id</i>	Determine the transaction ID	2 bytes	transaction ID

TABLE 13.21 Transaction information request item

isc_transaction_info() uses two buffers defined in the calling program: the *item-list buffer*, which lists transaction items about which you want information, and a *result buffer*, where the information requested is reported.

To define the *item-list buffer*, include the parameters *item_list_buffer_length* and *item_list_buffer_address*. The item-list buffer is a regular byte vector with no structure.

To define the result buffer, include the parameters *result_buffer_length* and *result_buffer_address*. These parameters specify the length and address of a buffer where the InterBase engine will place the return values from the function call.

The values returned to the result buffer are unaligned clusters of generic binary numbers. Furthermore, all numbers are represented in a generic format, with the least significant byte first, and the most significant byte last. Signed numbers have the sign in the last byte. Convert the numbers to a datatype native to your system before interpreting them.

In your call, include the item specifying the transaction ID, *isc_info_tra_id*. InterBase returns the transaction ID in the result buffer. In addition to the information InterBase returns in response to a request, InterBase can also return one or more of the following status messages to the result buffer. Each status message is one unsigned byte in length:

Item	Description
<i>isc_info_end</i>	End of the messages
<i>isc_info_truncated</i>	Result buffer is too small to hold any more requested information
<i>isc_info_error</i>	Requested information is unavailable; check the status vector for an error code and message

TABLE 13.22 Status message return items

The function return value indicates only that InterBase accepted the request for information. It does not mean that it understood the request or that it supplied all of the requested information. Your application must interpret the contents of the result buffer for details about the transaction.

Example The following code fragment gets information about a transaction:

```
static char          /* Declare item-list buffer. */
tra_items[] =
    {isc_info_tra_id};
/* Declare result buffer. */
CHAR tra_info[32];

isc_transaction_info(status_vector,
    &tr_handle,
    sizeof (tra_items), /* Length of item-list buffer. */
    &tra_items, /* Address of item-list buffer. */
    sizeof (tra_info), /* Length of result buffer. */
```

isc_vax_integer()

```
        &tra_info); /* Address of result buffer. */
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector);
    return(1);
}
```

Return Value `isc_transaction_info()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_start_transaction()`**

isc_vax_integer()

Reverses the byte order of an integer.

Syntax `ISC_LONG isc_vax_integer(
 char *buffer,
 short length);`

Parameter	Type	Description
<i>buffer</i>	<i>char*</i>	Pointer to the integer to convert
<i>length</i>	<i>short</i>	Length, in bytes, of the integer to convert Valid lengths are 1, 2, and 4 bytes

Description `isc_vax_integer()` reverses the byte order of an integer, specified in *buffer*, and returns the newly ordered value.

A typical use for this function is to convert integer values passed into a database parameter buffer to a format where the least significant byte must be first and the most significant byte last. In InterBase, integer values must be represented in input parameter buffers (for example, the DPB) and are returned in result buffers in a generic format where the least significant byte is first, and the most significant byte last.

`isc_vax_integer()` is used to convert integers to and from this format.

Example The following code fragment converts a 2-byte value, stored in a character buffer that is the result buffer returned by a function such as `isc_database_info()`:

```
#include <ibase.h>
char *p;
. . .
for(p = res_buffer; *p != isc_info_end;)
{
    /* Read item type of next cluster in the result buffer. */
    item = *p++;
    /* Read length of next value in result buffer, and convert. */
    len = isc_vax_integer(p, 2);
    p += len;
    /* Now process the actual value, len bytes in size. */
    . . .
}
```

Return Value `isc_vax_integer()` always returns a byte-reversed long integer value.

See Also [isc_attach_database\(\)](#), [isc_database_info\(\)](#)

isc_version()

Returns database implementation and version information.

Syntax

```
int isc_version(
    isc_db_handle *db_handle,
    isc_callback function_name,
    void *user_arg);
```

Parameter	Type	Description
<i>db_handle</i>	<i>isc_db_handle</i> *	<ul style="list-style-type: none"> • Pointer to a database handle set by a previous call to <code>isc_attach_database()</code> • <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>function_name</i>	<i>isc_callback</i>	<ul style="list-style-type: none"> • Pointer to a function to call with the relevant information • Passing a NULL pointer in C programs calls <code>printf()</code>
<i>user_arg</i>	<i>void</i> *	An application-specified parameter to pass as the first of two arguments to <i>function_name</i>

Description `isc_version()` determines the database implementation and on-disk structure (ODS) version numbers for the database specified by `db_handle`. It passes this information in two separate calls to the callback function pointed to by `function_name`.

`function_name` should point to an application function that takes two arguments: a void pointer, `user_arg`, and a char pointer. Applications can pass any kind of parameter desired in `user_arg`.

`isc_version()` makes two calls to `function_name`. First it determines the database implementation number, builds a string containing the information, and calls `function_name` with `user_arg`, and a pointer to the string containing the implementation number in the following format:

```
<implementation>(<class>), version "<version>"
```

where:

- *implementation* is a text string, such as “InterBase/NT”.
- *class* is a text string specifying the implementation class, such as “access method”.
- *version* is a version identification string, such as “4.0”.

The callback function specified by `function_name` is free to do with this information what it pleases.

After the callback function returns control to `isc_version()`, `isc_version()` builds a new string containing the ODS major and minor version numbers, then calls `function_name` a second time with `user_arg`, and a pointer to the string containing the ODS version number in the following format:

```
on disk structure version <ods_major_num>.<ods_minor_num>
```

where:

- *ods_major_num* is the major ODS number. Databases with different major version numbers have different physical layouts on disk and are incompatible with one another. A database engine can only access databases with a particular ODS major number.
- *ods_minor_num* is the minor ODS number. Differences in the minor ODS number, but not the major one indicate a non-structural change that still permits access by any database engine that recognizes the major version number.

Tip If a NULL pointer is passed for `function_name`, `isc_version()` sets `function_name` to point to the C `printf()` function.

Examples The following code fragment calls `isc_version()` with a NULL callback function:

```
#include <ibase.h>
. . .
int ret;
```

```

. . .
ret = isc_version(&db1, NULL, "\t%s\n");

```

Return Value If successful, `isc_version()` returns 0. Otherwise, it returns a nonzero value.

See Also [isc_database_info\(\)](#)

isc_wait_for_event()

Waits synchronously until one of a specified group of events is posted.

Note The `isc_wait_for_event()` function was called `gds_$event_wait()` in InterBase 3.3. It is therefore the only function that can't be translated from 3.3 nomenclature to all later versions by replacing `gds_$` with `isc_`.

Syntax

```

ISC_STATUS isc_wait_for_event(
    ISC_STATUS *status_vector,
    isc_db_handle *db_handle,
    short length,
    char *event_buffer,
    char *result_buffer);

```

Parameter	Type	Description
<i>status_vector</i>	<i>ISC_STATUS *</i>	Pointer to the error status vector
<i>db_handle</i>	<i>isc_db_handle *</i>	<ul style="list-style-type: none"> • Pointer to a database handle set by a previous call to <code>isc_attach_database()</code>; the handle identifies the database against which the events are expected to be posted • <i>db_handle</i> returns an error in <i>status_vector</i> if it is NULL
<i>length</i>	<i>short</i>	Length of the event parameter buffers, returned by the <code>isc_event_block()</code> call which allocated them
<i>event_buffer</i>	<i>char *</i>	Pointer to the event parameter buffer that specifies the current counts of the events to be waited on; this buffer should have been initially allocated and filled in by a call to <code>isc_event_block()</code>
<i>result_buffer</i>	<i>char *</i>	Pointer to the event parameter buffer to be filled in with updated event counts as a result of this function call; this buffer should have been initially allocated by a call to <code>isc_event_block()</code>

isc_wait_for_event()

Description `isc_wait_for_event()` is used to wait synchronously until one of a specified group of events is posted. Control is not returned to the calling application until one of the specified events occurs.

Events to wait on are specified in *event_buffer*, which should have been initially allocated and filled in by a previous call to `isc_event_block()`.

When one of these events is posted, `isc_wait_for_event()` fills in *result_buffer* with data that exactly corresponds to that in the initial buffer, except that the event counts will be the updated ones. Control then returns from `isc_wait_for_event()` to the calling application. The application should then call `isc_event_counts()` to determine which event was posted.

Note To request *asynchronous* notification of event postings, use `isc_que_events()` instead of `isc_wait_for_event()`. You *must* use asynchronous notifications in Microsoft Windows applications, or wherever a process must not stop processing.

Example The following program fragment illustrates a call to `isc_wait_for_event()` to wait for a posting of any of the events named “DEC”, “HP”, or “SUN”.

```
#include <ibase.h>
#define number_of_stocks 3;

char *event_buffer, *result_buffer;
short length;

length = (short)isc_event_block(
    &event_buffer,
    &result_buffer,
    number_of_stocks,
    "DEC", "HP", "SUN");

isc_wait_for_event(
    status_vector,
    &database_handle,
    length, /* Returned from isc_event_block(). */
    event_buffer,
    result_buffer);
if (status_vector[0] == 1 && status_vector[1])
{
    isc_print_status(status_vector); /* Display error message. */
    return(1);
}
```

```
/* Call isc_event_counts() to compare event counts in the buffers and  
thus determine which event(s) were posted. */
```

Return Value `isc_wait_for_event()` returns the second element of the status vector. Zero indicates success. A nonzero value indicates an error. For InterBase errors, the first element of the status vector is set to 1, and the second element is set to an InterBase error code.

To check for an InterBase error, examine the first two elements of the status vector directly. For more information about examining the status vector, see **Chapter 10, “Handling Error Conditions.”**

See Also **`isc_event_block()`, `isc_que_events()`**

A

InterBase Document Conventions

This appendix covers the following topics:

- The InterBase 6 documentation set
- The printing conventions used to display information in text
- The printing conventions used to display information in syntax, code, and examples

The InterBase documentation set

The InterBase documentation set is an integrated package designed for all levels of users. It consists of six full-length printed books plus the *Installation Guide*. Each of these books is also provided in Adobe Acrobat PDF format and is accessible on line. If Adobe Acrobat is not already installed on your system, you can find it on the InterBase distribution CD-ROM or at <http://www.adobe.com/prodindex/acrobat/readstep.html>. Acrobat is available for Windows NT, Windows 95, and most flavors of UNIX.

Book	Description
<i>Operations Guide</i>	Provides an introduction to InterBase and an explanation of tools and procedures for performing administrative tasks on databases and database servers; also includes full reference on InterBase utilities, including isql , gbak , gfix , and others
<i>Data Definition Guide</i>	Explains how to create, alter, and delete database objects using the SQL language
<i>Developer's Guide</i>	Provides both reference and task-oriented material for users of the Borland RAD tools (Delphi, C++ Builder, and JBuilder); includes chapters on writing UDFs, driver configuration, developing embedded installation applications, and using the new InterBase Data Access Components
<i>Language Reference</i>	Describes the SQL language syntax and usage; includes references for procedure and trigger language, InterBase keywords, functions in the InterBase UDF library, error codes, character sets, and the system tables
<i>Embedded SQL Guide</i>	(formerly called the <i>Programmer's Guide</i>) Describes how to write embedded SQL database applications in a host language, precompiled through gpre
<i>API Guide</i>	Explains how to write database applications using the InterBase API

TABLE A.1 Books in the InterBase 6 documentation set

Printing conventions

The InterBase documentation set uses various typographic conventions to identify objects and syntactic elements.

The following table lists typographic conventions used in text, and provides examples of their use:

Convention	Purpose	Example
UPPERCASE	SQL keywords, SQL functions, and names of all database objects such as tables, columns, indexes, and stored procedures	<ul style="list-style-type: none"> the SELECT statement retrieves data from the CITY column in the CITIES table can be used in CHAR, VARCHAR, and BLOB text columns the CAST() function
<i>italic</i>	New terms, emphasized words, all elements from host languages, and all user-supplied items	<ul style="list-style-type: none"> <i>isc_decode_date()</i> the host variable, <i>segment_length</i> contains six variables, or <i>data members</i>
bold	File names, menu picks, and all commands that are entered at a system prompt, including their switches, arguments, and parameters	<ul style="list-style-type: none"> gbak, isql, gsec, gfix specify the gpre -sqlda old switch a script, ib_udf.sql, in the examples subdirectory the employee.gdb database; the employee database the Session Advanced Settings command

TABLE A.2 Text conventions

Syntax conventions

The following table lists the conventions used in syntax statements and sample code, and provides examples of their use:

Convention	Purpose	Example
UPPERCASE	Keywords that must be typed exactly as they appear when used	<ul style="list-style-type: none"> •SET TERM !!; •ADD [CONSTRAINT] CHECK
<i>italic</i>	User-supplied parameters that cannot be broken into smaller units	<ul style="list-style-type: none"> •CREATE TRIGGER <i>name</i> FOR <i>table</i>; •ALTER EXCEPTION <i>name</i> '<i>message</i>'
< <i>italic</i> >	Parameters in angle brackets can be broken into smaller syntactic units; the expansion syntax for these parameters follows the current syntax statement	WHILE (<condition>) DO <compound_statement>
[]	Optional syntax: you do not need to include anything that is enclosed in square brackets; when elements within these brackets are separated by the pipe symbol (), you can choose only one	<ul style="list-style-type: none"> •CREATE [UNIQUE][ASCENDING DESCENDING] •[FILTER [FROM <i>subtype</i>] TO <i>subtype</i>]
{ }	You must include one and only one of the enclosed options, which are separated by the pipe symbol ()	{INTO USING}
	You can choose only one of a group whose elements are separated by this pipe symbol	SELECT [DISTINCT ALL]
...	You can repeat the clause enclosed in brackets with the "... " symbol as many times as necessary	(<col> [, <col>...])

TABLE A.3 Syntax conventions

B

Data Structures

This appendix documents the data structures, compile-time constants, parameter buffers, and information buffers utilized in InterBase API applications.

This information also appears throughout the rest of this API Guide, but is consolidated here as a convenience. See other sections of this manual for more information.

All the structures and compile-time constants mentioned are defined in the **ibase.h** header file. Items are documented alphabetically, as follows:

- Array descriptor
- Blob descriptor
- Blob information item-list buffer and result buffer
- Blob parameter buffer
- Database information request buffer and result buffer
- Database parameter buffer
- SQL datatype macro constants
- Status vector
- Transaction parameter buffer
- *XSQLDA* and *XSQLVAR* structures

Array descriptor

An array descriptor `ISC_ARRAY_DESC` is a structure defined in the `ibase.h` header file as:

```
typedef struct {
    unsigned char array_desc_dtype;
    char array_desc_scale;
    unsigned short array_desc_length;
    char array_desc_field_name [32];
    char array_desc_relation_name [32];
    short array_desc_dimensions;
    short array_desc_flags;
    ISC_ARRAY_BOUND array_desc_bounds [16];
} ISC_ARRAY_DESC;
```

`ISC_ARRAY_BOUND` is defined as:

```
typedef struct {
    short array_bound_lower; /* lower bound */
    short array_bound_upper; /* upper bound */
} ISC_ARRAY_BOUND;
```

Field	Description
<i>array_desc_dtype</i>	Datatype (see below)
<i>array_desc_scale</i>	Scale for numeric datatypes
<i>array_desc_length</i>	Length in bytes of each array element
<i>array_desc_field_name</i>	NULL-terminated column name
<i>array_desc_relation_name</i>	NULL-terminated relation name
<i>array_desc_dimensions</i>	Number of array dimensions
<i>array_desc_flags</i>	Specifies whether array is to be accessed in row- major or column-major order <ul style="list-style-type: none"> • 0: row-major • 1: column-major
<i>array_desc_bounds</i>	Lower and upper bounds for each dimension

TABLE B.1 Array descriptor fields

Datatypes for array descriptors

The *array_desc_dtype* field of an array descriptor must be expressed as one of the datatypes in the following table:

array_desc_dtype value	Corresponding InterBase datatype
<i>blr_text</i>	CHAR
<i>blr_text2</i>	CHAR
<i>blr_short</i>	SMALLINT
<i>blr_long</i>	INTEGER
<i>blr_quad</i>	ISC_QUAD structure
<i>blr_float</i>	FLOAT
<i>blr_double</i>	DOUBLE PRECISION
<i>blr_sql_date</i>	DATE
<i>blr_sql_time</i>	TIME
<i>blr_timestamp</i>	TIMESTAMP
<i>blr_varying</i>	VARCHAR
<i>blr_varying2</i>	VARCHAR
<i>blr_blob_id</i>	ISC_QUAD structure
<i>blr_cstring</i>	NULL-terminated string
<i>blr_cstring2</i>	NULL-terminated string

TABLE B.2 Datatypes for array descriptors

Blob descriptor

A Blob descriptor is defined as:

```
typedef struct {
    short blob_desc_subtype;
    short blob_desc_charset;
```

```

short blob_desc_segment_size;
unsigned char blob_desc_field_name [32];
unsigned char blob_desc_relation_name [32];
} ISC_BLOB_DESC;

```

Field	Description
<i>blob_desc_subtype</i>	Type of Blob data 0: unstructured 1: TEXT negative integer between -1 and -32678: user-defined subtype
<i>blob_desc_charset</i>	Character set (see below)
<i>blob_desc_segment_size</i>	Segment size
<i>blob_desc_field_name</i>	NULL-terminated column name
<i>blob_desc_relation_name</i>	NULL-terminated table name

TABLE B.3 Blob descriptor fields

Character sets

InterBase supports a number of character sets. For a list of the character sets supported, and the character set value that must be entered in the *blob_desc_charset* field of a Blob descriptor, see the *Data Definition Guide*.

Blob information buffers

The *isc_blob_info()* call enables an application to query for Blob information such as the total number of segments in the Blob, or the length of the longest segment.

isc_blob_info() requires two application-provided buffers, an item-list buffer, where the application specifies the information it needs, and a result buffer, where InterBase returns the requested information. An application populates the item-list buffer with information prior to calling *isc_blob_info()*. InterBase returns information in the result buffer. If InterBase attempts to pass back more information than can fit in the result buffer, it puts the value, *isc_info_truncated*, defined in *ibase.h*, in the final byte of the result buffer.

► *Item-list buffer*

The item-list buffer is a byte vector into which is placed a sequence of byte values, one per requested item of information. Each byte is an *item type*, specifying the kind of information desired. Compile-time constants for all item types are defined in **ibase.h**:

```
#define isc_info_blob_num_segments      4
#define isc_info_blob_max_segment      5
#define isc_info_blob_total_length     6
#define isc_info_blob_type             7
```

► *Result buffer*

The result buffer is a byte vector in which InterBase returns a series of clusters of information, one per item requested. Each cluster consists of three parts:

1. A one-byte *item type*. Each is the same as one of the item types in the item-list buffer.
2. A 2-byte number specifying the number of bytes that follow in the remainder of the cluster.
3. A *value*, stored in a variable number of bytes, whose interpretation depends on the item type.

A calling program is responsible for interpreting the contents of the result buffer and for deciphering each cluster as appropriate.

The clusters returned to the result buffer are not aligned. Furthermore, all numbers are represented in a generic format, with the least significant byte first, and the most significant byte last. Signed numbers have the sign in the last byte. Convert the numbers to a datatype native to your system, if necessary, before interpreting them. The API call, *isc_vax_integer()*, can be used to perform the conversion.

Blob buffer items

The following table lists items about which information can be requested and returned, and the values reported:

Request and return item	Return value
<i>isc_info_blob_num_segments</i>	Total number of segments
<i>isc_info_blob_max_segment</i>	Length of the longest segment
<i>isc_info_blob_total_length</i>	Total size, in bytes, of Blob
<i>isc_info_blob_type</i>	Type of Blob (0: segmented, or 1: stream)

TABLE B.4 Blob information items and return values

In addition to the information InterBase returns in response to a request, InterBase can also return one or more of the following status messages to the result buffer. Each status message is one unsigned byte in length:

Item	Description
<i>isc_info_end</i>	End of the messages
<i>isc_info_truncated</i>	Result buffer is too small to hold any more requested information
<i>isc_info_error</i>	Requested information is unavailable. Check the status vector for an error code and message

TABLE B.5 Status message return items

Blob parameter buffer

A *Blob Parameter Buffer* (BPB) is an application-defined byte vector, passed as an argument to *isc_open_blob2()* or *isc_create_blob2()*, that specifies Blob attributes required for filtering Blob data.

A BPB consists of the following parts:

1. A byte specifying the version of the parameter buffer, always the compile-time constant, *isc_bpb_version1*.
2. A contiguous series of one or more clusters of bytes, each describing a single parameter.

Each cluster consists of the following parts:

1. A one-byte parameter type. There are compile-time constants defined for all the parameter types (for example, *isc_bpb_target_type*).
2. A one-byte number specifying the number of bytes that follow in the remainder of the cluster.
3. A variable number of bytes, whose interpretation depends on the parameter type.

All numbers in the Blob parameter buffer must be represented in a generic format, with the least significant byte first, and the most significant byte last. Signed numbers should have the sign in the last byte. The API function *isc_vax_integer()* can be used to reverse the byte order of a number. For more information about *isc_vax_integer()*, see **“isc_vax_integer()” on page 391 of Chapter 13, “API Function Reference.”**

The following table lists the parameter types and their meaning. For lists of the possible subtypes and character sets, see “Blob descriptor” on page 403.

Parameter type	Description
<i>isc_bpb_target_type</i>	Target subtype
<i>isc_bpb_source_type</i>	Source subtype
<i>isc_bpb_target_interp</i>	Target character set
<i>isc_bpb_source_interp</i>	Source character set

TABLE B.6 Blob parameter buffer parameter types

Database information request buffer and result buffer

The *isc_database_info()* call enables an application to query for information about an attached database.

isc_database_info() requires two application-provided buffers, a request buffer, where the application specifies the information it needs, and a result buffer, where InterBase returns the requested information. An application populates the request buffer with information prior to calling *isc_database_info()*. InterBase returns information in the result buffer. If InterBase attempts to pass back more information than can fit in the result buffer, it puts the value, *isc_info_truncated*, defined in *ibase.h*, in the final byte of the result buffer.

Request buffer

The request buffer is a byte vector into which is placed a sequence of byte values, one per requested item of information. Each byte is an *item type*, specifying the kind of information desired. Compile-time constants for all item types are defined in **ibase.h** and shown below.

Result buffer

The result buffer is a byte vector in which InterBase returns a series of clusters of information, one per item requested. Each cluster consists of three parts:

1. A one-byte *item return type*. These are the same as the item types specified in the request buffer.
2. A two-byte number specifying the number of bytes that follow in the remainder of the cluster.
3. A *value*, stored in a variable number of bytes, whose interpretation (as a number or as a string of characters, for example) depends on the item return type.

A calling program is responsible for interpreting the contents of the result buffer and for deciphering each cluster as appropriate. In many cases, the value simply contains a number or a string (sequence of characters). But in other cases, the value is a number of bytes whose interpretation depends on the item return type.

The clusters returned to the result buffer are not aligned. Furthermore, all numbers are represented in a generic format, with the least significant byte first, and the most significant byte last. Signed numbers have the sign in the last byte. Convert the numbers to a datatype native to your system, if necessary, before interpreting them. The API call, *isc_vax_integer()*, can be used to perform the conversion.

In addition to the information InterBase returns in response to a request, InterBase can also return one or more of the following status messages to the result buffer. Each status message is one unsigned byte in length:

Item	Description
<i>isc_info_end</i>	End of the messages
<i>isc_info_truncated</i>	Result buffer is too small to hold any more requested information
<i>isc_info_error</i>	Requested information is unavailable; check the status vector for an error code and message

TABLE B.7 Status message return items

Request buffer items and result buffer values

The following sections show the request buffer items and result buffer contents for the following categories of database information:

- Database characteristics
- Environmental characteristics
- Performance statistics
- Database operation counts

► *Database characteristics*

Several items are provided for determining database characteristics, such as its size and major and minor ODS version numbers. The following table lists the request buffer items that can be passed, and the information returned in the result buffer for each item type:

Request buffer item	Result buffer contents
<i>isc_info_allocation</i>	Number of database pages allocated
<i>isc_info_base_level</i>	Database version (level) number: 1 byte containing the number 1 1 byte containing the version number
<i>isc_info_db_id</i>	Database file name and site name: • 1 byte containing the number 2 • 1 byte containing the length, <i>d</i> , of the database file name in bytes • A string of <i>d</i> bytes, containing the database file name • 1 byte containing the length, <i>l</i> , of the site name in bytes • A string of <i>l</i> bytes, containing the site name
<i>isc_info_implementation</i>	Database implementation number: • 1 byte containing a 1 • 1 byte containing the implementation number • 1 byte containing a "class" number, either 1 or 12

TABLE B.8 Database information items for database characteristics

Request buffer item	Result buffer contents
<i>isc_info_no_reserve</i>	0 or 1 <ul style="list-style-type: none"> • 0 indicates space is reserved on each database page for holding backup versions of modified records [Default] • 1 indicates no space is reserved for such records
<i>isc_info_ods_minor_version</i>	On-disk structure (ODS) minor version number; an increase in a minor version number indicates a non-structural change, one that still allows the database to be accessed by database engines with the same major version number but possibly different minor version numbers
<i>isc_info_ods_version</i>	ODS major version number <ul style="list-style-type: none"> • Databases with different major version numbers have different physical layouts; a database engine can only access databases with a particular ODS major version number • Trying to attach to a database with a different ODS number results in an error
<i>isc_info_page_size</i>	Number of bytes per page of the attached database; use with <i>isc_info_allocation</i> to determine the size of the database
<i>isc_info_version</i>	Version identification string of the database implementation: <ul style="list-style-type: none"> • 1 byte containing the number 1 • 1 byte specifying the length, <i>n</i>, of the following string • <i>n</i> bytes containing the version identification string

TABLE B.8 Database information items for database characteristics (continued)

► *Environmental characteristics*

Several items are provided for determining environmental characteristics, such as the amount of memory currently in use, or the number of database cache buffers currently allocated. These items are described in the following table:

Request buffer item	Result buffer contents
<i>isc_info_current_memory</i>	Amount of server memory (in bytes) currently in use
<i>isc_info_forced_writes</i>	Number specifying the mode in which database writes are performed (0 for asynchronous, 1 for synchronous)
<i>isc_info_max_memory</i>	Maximum amount of memory (in bytes) used at one time since the first process attached to the database
<i>isc_info_num_buffers</i>	Number of memory buffers currently allocated
<i>isc_info_sweep_interval</i>	Number of transactions that are committed between “sweeps” to remove database record versions that are no longer needed
<i>isc_info_user_names</i>	NetWare only. Names of all the users currently attached to the database; for <i>each</i> such user, the result buffer will contain an <i>isc_info_user_names</i> byte followed by a 1-byte length specifying the number of bytes in the user name, followed by the user name

TABLE B.9 Database information items for environmental characteristics

Not all environmental information items are available on all platforms.

► *Performance statistics*

There are four items providing performance statistics for a database. These statistics accumulate for a database from the moment it is first attached by any process until the last remaining process detaches from the database. A program requesting this information, therefore, sees information pertaining to its own attachment and all other attachments.

For example, the value returned for *isc_info_reads* is the number of reads since the current database was first attached: it is an *aggregate* of all reads done by all attached processes, rather than the number of reads done for the calling program since it attached to the database.

The items providing performance statistics are summarized in the following table:

Request buffer item	Result buffer contents
<i>isc_info_fetches</i>	Number of reads from the memory buffer cache
<i>isc_info_marks</i>	Number of writes to the memory buffer cache
<i>isc_info_reads</i>	Number of page reads
<i>isc_info_writes</i>	Number of page writes

TABLE B.10 Database information items for performance statistics

► *Database operation counts*

Several information items are provided for determining the number of various database operations performed by the currently attached calling program. These values are calculated on a per-table basis.

When any of these information items is requested, InterBase returns to the result buffer:

1. 1 byte specifying the item type (for example, *isc_info_insert_count*).
2. 2 bytes telling how many bytes compose the subsequent value pairs.
3. A pair of values for each table in the database on which the requested type of operation has occurred since the database was last attached.

Each pair consists of:

- 2 bytes specifying the table ID.
- 4 bytes listing the number of operations (for example, inserts) done on that table.

To determine an actual table name from a table ID, query the RDB\$RELATION system table.

The following table describes the items which return count values for operations on the database:

Request buffer item	Result buffer contents
<i>isc_info_backout_count</i>	Number of removals of a version of a record
<i>isc_info_delete_count</i>	Number of database deletes since the database was last attached
<i>isc_info_expunge_count</i>	Number of removals of a record and all of its ancestors, for records whose deletions have been committed
<i>insert_count</i>	Number of inserts into the database since the database was last attached
<i>purge_count</i>	Number of removals of old versions of fully mature records (records that are committed, so that older ancestor versions are no longer needed)
<i>read_idx_count</i>	Number of reads done via an index since the database was last attached
<i>read_seq_count</i>	Number of sequential sequential table scans (row reads) done on each table since the database was last attached
<i>update_count</i>	Number of database updates since the database was last attached

TABLE B.11 Database information items for operation counts

A *Database Parameter Buffer* (DPB) is an application-defined byte vector, passed as an argument to *isc_attach_database()*, that specifies desired database characteristics.

A DPB consists of the following parts:

1. A byte specifying the version of the parameter buffer, always the compile-time constant, *isc_dpb_version1*.
2. A contiguous series of one or more clusters of bytes, each describing a single parameter.

Each cluster consists of the following parts:

1. A one-byte parameter type. There are compile-time constants defined for all the parameter types (for example, *isc_dpb_num_buffers*).
2. A one-byte number specifying the number of bytes that follow in the remainder of the cluster.

3. A variable number of bytes, whose interpretation (as a number or as a string of characters, for example) depends on the parameter type.

The following table lists DPB items by purpose:

User Validation	Item
User name	<i>isc_dpb_user_name</i>
Password	<i>isc_dpb_password</i>
Encrypted password	<i>isc_dpb_password_enc</i>
System database administrator's user name	<i>isc_dpb_sys_user_name</i>
Authorization key for a software license	<i>isc_dpb_license</i>
Database encryption key	<i>isc_dpb_encrypt_key</i>
Environmental control	
Number of cache buffers	<i>isc_dpb_num_buffers</i>
<i>dbkey</i> context scope	<i>isc_dpb_dbkey_scope</i>
System management	
Force writes to the database to be done asynchronously or synchronously	<i>isc_dpb_force_write</i>
Specify whether or not to reserve a small amount of space on each database page for holding backup versions of records when modifications are made.	<i>isc_dpb_no_reserve</i>
Specify whether or not the database should be marked as damaged	<i>isc_dpb_damaged</i>
Perform consistency checking of internal structures	<i>isc_dpb_verify</i>
Shadow control	
Activate the database shadow, an optional, duplicate, in-sync copy of the database	<i>isc_dpb_activate_shadow</i>
Delete the database shadow	<i>isc_dpb_delete_shadow</i>

TABLE B.12 DPB parameters

User Validation	Item
Replay logging system control	
Activate a replay logging system to keep track of all database calls	<i>isc_dpb_begin_log</i>
Deactivate the replay logging system	<i>isc_dpb_quit_log</i>
Character set and message file specification	
Language-specific message file	<i>isc_dpb_lc_messages</i>
Character set to be utilized	<i>isc_dpb_lc_ctype</i>

TABLE B.12 DPB parameters (*continued*)

The following table lists DPB parameters in alphabetical order. For each parameter, it lists its purpose, the length, in bytes, of any values passed with the parameter, and the value to pass:

Parameter	Purpose	Length	Value
<i>isc_dpb_activate_shadow</i>	Directive to activate the database shadow, which is an optional, duplicate, in-sync copy of the database	1 (Ignored)	0 (Ignored)
<i>isc_dpb_damaged</i>	Number signifying whether or not the database should be marked as damaged; 1 = mark as damaged, 0 = do <i>not</i> mark as damaged	1	0 or 1
<i>isc_dpb_dbkey_scope</i>	Scope of <i>dbkey</i> context; 0 limits scope to the current transaction, 1 extends scope to the database session.	1	0 or 1
<i>isc_dpb_delete_shadow</i>	Directive to delete a database shadow that is no longer needed	1 (Ignored)	0 (Ignored)
<i>isc_dpb_encrypt_key</i>	String encryption key, up to 255 characters	Number of bytes in string	String containing key
<i>isc_dpb_force_write</i>	Specifies whether database writes are synchronous or asynchronous; 0 = asynchronous, 1 = synchronous	1	0 or 1

TABLE B.13 Alphabetical list of DPB parameters

Parameter	Purpose	Length	Value
<i>isc_dpb_lc_ctype</i>	String specifying the character set to be utilized	Number of bytes in string	String containing character set name
<i>isc_dpb_lc_messages</i>	String specifying a language-specific message file	Number of bytes in string	String containing message file name
<i>isc_dpb_license</i>	String authorization key for a software license	Number of bytes in string	String containing key
<i>isc_dpb_no_reserve</i>	Specifies whether or not a small amount of space on each database page is reserved for holding backup versions of records when modifications are made; keep backup versions on the same page as the primary record to optimize update activity 0 (default) = reserve space 1= do not reserve space	1	0 or 1
<i>isc_dpb_num_buffers</i>	Number of database cache buffers to allocate for use with the database Default = 75	1	Number of buffers to allocate
<i>isc_dpb_password</i>	String password, up to 255 characters	Number of bytes in string	String containing password
<i>isc_dpb_password_enc</i>	String encrypted password, up to 255 characters	Number of bytes in string	String containing password
<i>isc_dpb_sys_user_name</i>	String system DBA name, up to 255 characters	Number of bytes in string	String containing SYSDBA name
<i>isc_dpb_user_name</i>	String user name, up to 255 characters	Number of bytes in string	String containing user name

TABLE B.13 Alphabetical list of DPB parameters (continued)

Some parameters, such as *isc_dpb_delete_shadow*, are directives that do not require additional parameters. Even so, you *must* still provide length and value bytes for these parameters. Set length to 1, and value to 0. InterBase ignores these parameter values, but they are required to maintain the format of the DPB.

SQL datatype macro constants

InterBase defines a set of macro constants to represent SQL datatypes and NULL status information in an *XSQLVAR*. An application should use these macro constants to specify the datatype of parameters and to determine the datatypes of select-list items in an SQL statement. The following table lists each SQL datatype, its corresponding macro constant expression, C datatype or InterBase typedef, and whether or not the *sqlind* field is used to indicate a parameter or variable that contains NULL or unknown data:

SQL datatype	Macro expression	C datatype or typedef	<i>sqlind</i> used?
Array	SQL_ARRAY	ISC_QUAD	No
Array	SQL_ARRAY + 1	ISC_QUAD	Yes
Blob	SQL_BLOB	ISC_QUAD	No
Blob	SQL_BLOB + 1	ISC_QUAD	Yes
CHAR	SQL_TEXT	char[]	No
CHAR	SQL_TEXT + 1	char[]	Yes
DATE	SQL_DATE	ISC_DATE	No
DATE	SQL_DATE + 1	ISC_DATE	Yes
DECIMAL	SQL_SHORT, SQL_LONG, SQL_DOUBLE, or SQL_INT64	int, long, double, or ISC_INT64	No
DECIMAL	SQL_SHORT + 1, SQL_LONG + 1, SQL_DOUBLE + 1, or SQL_INT64 + 1	int, long, double, or ISC_INT64	Yes
DOUBLE PRECISION	SQL_DOUBLE	double	No
DOUBLE PRECISION	SQL_DOUBLE + 1	double	Yes
INTEGER	SQL_LONG	long	No
INTEGER	SQL_LONG + 1	long	Yes
FLOAT	SQL_FLOAT	float	No
FLOAT	SQL_FLOAT + 1	float	Yes

TABLE B.14 SQL datatypes, macro expressions, and C datatypes

SQL datatype	Macro expression	C datatype or typedef	<i>sqlind</i> used?
NUMERIC	SQL_SHORT, SQL_LONG, SQL_DOUBLE, or SQL_INT64	int, long, double, or ISC_INT64	No
NUMERIC	SQL_SHORT + 1, SQL_LONG + 1, SQL_DOUBLE + 1, or SQL_INT64 + 1	int, long, double, or ISC_INT64	Yes
SMALLINT	SQL_SHORT	short	No
SMALLINT	SQL_SHORT + 1	short	Yes
TIME	SQL_TIME	ISC_TIME	No
TIME	SQL_TIME + 1	ISC_TIME	Yes
TIMESTAMP	SQL_TIMESTAMP	ISC_TIMESTAMP	No
TIMESTAMP	SQL_TIMESTAMP + 1	ISC_TIMESTAMP	Yes
VARCHAR	SQL_VARYING	First 2 bytes: short containing the length of the character string. Remaining bytes: char[]	No
VARCHAR	SQL_VARYING + 1	First 2 bytes: short containing the length of the character string. Remaining bytes: char[]	Yes

TABLE B.14 SQL datatypes, macro expressions, and C datatypes (*continued*)

DECIMAL and NUMERIC datatypes are stored internally as SMALLINT, INTEGER, DOUBLE PRECISION, or 64-bit integer datatypes. To specify the correct macro expression to provide for a DECIMAL or NUMERIC column, use **isql** to examine the column definition in the table to see how InterBase is storing column data, then choose a corresponding macro expression.

The datatype information for a parameter or select-list item is contained in the *sqltype* field of the *XSQLVAR* structure. The value contained in *sqltype* provides two pieces of information:

- The datatype of the parameter or select-list item.
- Whether *sqlind* is used to indicate NULL values. If *sqlind* is used, its value specifies whether the parameter or select-list item is NULL (-1), or not NULL (0).

For example, if *sqltype* equals SQL_TEXT, the parameter or select-list item is a CHAR that does not use *sqlind* to check for a NULL value (because, in theory, NULL values are not allowed for it). If *sqltype* equals SQL_TEXT + 1, then *sqlind* can be checked to see if the parameter or select-list item is NULL.

The C language expression, *sqltype & 1*, provides a useful test of whether a parameter or select-list item can contain a NULL. The expression evaluates to 0 if the parameter or select-list item cannot contain a NULL, and 1 if the parameter or select-list item can contain a NULL.

Status vector

Most API functions return status information that indicates success or failure. Status information is reported in an error status vector, declared in applications as an array of twenty long integers, using the following syntax:

```
#include <ibase.h>
. . .
ISC_STATUS status_vector[20];
```

If you plan to write your own routines instead of the InterBase error-handling routines to read and react to the contents of the status vector, you need to know how to interpret it.

InterBase stores error information in the status vector in clusters of two or three longs. The first cluster in the status vector *always* indicates the primary cause of the error. Subsequent clusters may contain supporting information about the error, for example, strings or numbers for display in an associated error message. The actual number of clusters used to report supporting information varies from error to error.

In many cases, additional errors may be reported in the status vector. Additional errors are reported immediately following the first error and its supporting information, if any. The first cluster for each additional error message identifies the error. Subsequent clusters may contain supporting information about the error.

Meaning of the first long in a cluster

The first long in any cluster is a *numeric descriptor*. By examining the numeric descriptor for a cluster, you can always determine the:

- Total number of longs in the cluster.
- Kind of information reported in the remainder of the cluster.
- Starting location of the next cluster in the status vector.

The following table lists possible values for the first long in any cluster in the status vector. Note that the first cluster in the status vector can only contain values of 0, 1, or greater than 4:

Value	Longword in cluster	Meaning
0	—	End of error information in the status vector
1	2	Second long is an InterBase error code
2	2	Second long is the address of string used as a replaceable parameter in a generic InterBase error message
3	3	Second long is the length, in bytes, of a variable length string provided by the operating system (most often this string is a file name); third long is the address of the string
4	2	Second long is a number used as a replaceable parameter in a generic InterBase error message
5	2	Second long is the address of an error message string requiring no further processing before display
6	2	Second long is a VAX/VMS error code
7	2	Second long is a UNIX error code
8	2	Second long is an Apollo Domain error code
9	2	Second long is an MS-DOS or OS/2 error code
10	2	Second long is an HP MPE/XL error code
11	2	Second long is an HP MPE/XL IPC error code
12	2	Second long is a NeXT/Mach error code

NOTE As InterBase is adapted to run on other hardware and software platforms, additional numeric descriptors for specific platform and operating system error codes may be added to the end of this list.

TABLE B.15 Interpretation of status vector clusters

The following table lists the `ibase.h` `#define` equivalents of each numeric descriptor:

Value	#define	Value	#define
0	<code>isc_arg_end</code>	8	<code>isc_arg_domain</code>
1	<code>isc_arg_gds</code>	9	<code>isc_arg_dos</code>
2	<code>isc_arg_string</code>	10	<code>isc_arg_mpexl</code>
3	<code>isc_arg_cstring</code>	11	<code>isc_arg_mpexl_ipc</code>
4	<code>isc_arg_number</code>	15	<code>isc_arg_next_mach</code>
5	<code>isc_arg_interpreted</code>	16	<code>isc_arg_netware</code>
6	<code>isc_arg_vms</code>	17	<code>isc_arg_win32</code>
7	<code>isc_arg_unix</code>		

TABLE B.16 `#defines` for status vector numeric descriptors

Transaction parameter buffer

The *transaction parameter buffer* (TPB) is an optional, application-defined byte vector, passed as an argument to `isc_start_transaction()`, that sets up a transaction's *attributes*, its operating characteristics, such as whether the transaction has read and write access to tables, or read-only access, and whether or not other simultaneously active transactions can share table access with the transaction. Each transaction may have its own TPB, or transactions that share operating characteristics can use the same TPB.

If a TPB is not created for a transaction, a NULL pointer must be passed to `isc_start_transaction()` in its place. A default set of attributes is automatically assigned to such transactions.

A TPB is declared in a C program as a char array of one-byte elements. Each element is a parameter that describes a single transaction attribute. The first element in every TPB must be the *isc_tpb_version3* constant. The following table lists available TPB constants, describes their purposes, and indicates which constants are assigned as a default set of attributes when a NULL TPB pointer is passed to *isc_start_transaction()*:

Parameter	Description
<i>isc_tpb_version3</i>	InterBase version 3 transaction
<i>isc_tpb_consistency</i>	Table-locking transaction model
<i>isc_tpb_concurrency</i>	High throughput, high concurrency transaction with acceptable consistency; use of this parameter takes full advantage of the InterBase multi-generational transaction model [Default]
<i>isc_tpb_shared</i>	Concurrent, shared access of a specified table among all transactions.; use in conjunction with <i>isc_tpb_lock_read</i> and <i>isc_tpb_lock_write</i> to establish the lock option [Default]
<i>isc_tpb_protected</i>	Concurrent, restricted access of a specified table; use in conjunction with <i>isc_tpb_lock_read</i> and <i>isc_tpb_lock_write</i> to establish the lock option
<i>isc_tpb_wait</i>	Lock resolution specifies that the transaction is to wait until locked resources are released before retrying an operation [Default]
<i>isc_tpb_nowait</i>	Lock resolution specifies that the transaction is not to wait for locks to be released, but instead, a lock conflict error should be returned immediately
<i>isc_tpb_read</i>	Access mode of read-only that allows a transaction only to select data from tables
<i>isc_tpb_write</i>	Access mode of read-write that allows a transaction to select, insert, update, and delete table data [Default]
<i>isc_tpb_lock_read</i>	Read-only access of a specified table; use in conjunction with <i>isc_tpb_shared</i> , <i>isc_tpb_protected</i> , and <i>isc_tpb_exclusive</i> to establish the lock option
<i>isc_tpb_lock_write</i>	Read-write access of a specified table; use in conjunction with <i>isc_tpb_shared</i> , <i>isc_tpb_protected</i> , and <i>isc_tpb_exclusive</i> to establish the lock option [Default]

TABLE B.17 TPB constants

Parameter	Description
<i>isc_tpb_read_committed</i>	High throughput, high concurrency transaction that can read changes committed by other concurrent transactions; use of this parameter takes full advantage of the InterBase multi-generational transaction model
<i>isc_tpb_rec_version</i>	Enables an <i>isc_tpb_read_committed</i> transaction to read the most recently committed version of a record even if other, uncommitted versions are pending.
<i>isc_tpb_no_rec_version</i>	Enables an <i>isc_tpb_read_committed</i> transaction to read only the latest committed version of a record If an uncommitted version of a record is pending and <i>isc_tpb_wait</i> is also specified, then the transaction waits for the pending record to be committed or rolled back before proceeding; otherwise, a lock conflict error is reported at once

TABLE B.17 TPB constants (continued)

TPB parameters specify the following classes of information:

- *Transaction version number* is used internally by the InterBase engine. It is always be the first attribute specified in the TPB, and must always be set to *isc_tpb_version3*.
- *Access mode* describes the actions that can be performed by the functions associated with the transaction. Valid access modes are:

isc_tpb_read

isc_tpb_write

- *Isolation level* describes the view of the database given a transaction as it relates to actions performed by other simultaneously occurring transactions. Valid isolation levels are:

isc_tpb_concurrency

isc_tpb_consistency

isc_tpb_read_committed, isc_tpb_rec_version

isc_tpb_read_committed, isc_tpb_no_rec_version

- *Lock resolution* describes how a transaction should react if a lock conflict occurs. Valid lock resolutions are:

isc_tpb_wait

isc_tpb_nowait

- *Table reservation* optionally describes an access method and lock resolution for a specified table that the transaction accesses. When table reservation is used, tables are reserved for the specified access when the transaction is started, rather than when the transaction actually accesses the table. Valid reservations are:

isc_tpb_shared, isc_tpb_lock_write

isc_tpb_shared, isc_tpb_lock_read

isc_tpb_protected, isc_tpb_lock_write

isc_tpb_protected, isc_tpb_lock_read

XSQLDA and XSQLVAR

All DSQL applications must declare one or more extended SQL descriptor areas (*XSQLDAs*).


The *XSQLDA* is a host-language data structure that DSQL uses to transport data to or from a database when processing an SQL statement string. There are two types of *XSQLDAs*: *input* descriptors and *output* descriptors. Both input and output descriptors are implemented using the *XSQLDA* structure.

Syntax

One field in the *XSQLDA*, *sqlvar*, is an *XSQLVAR* structure. The *sqlvar* is especially important because one *XSQLVAR* must be defined for each input parameter or column returned.

Applications do not declare instances of the *XSQLVAR* ahead of time, but must, instead, dynamically allocate storage for the proper number of *XSQLVAR* structures required for each DSQL statement before it is executed, then deallocate it, as appropriate, after statement execution.

The following figure illustrates the relationship between the *XSQLDA* and the *XSQLVAR*.



An input *XSQLDA* consists of a single *XSQLDA* structure, and one *XSQLVAR* structure for each input parameter. An output *XSQLDA* also consists of one *XSQLDA* structure and one *XSQLVAR* structure for each data item returned by the statement.

The `isc_dsql_prepare()`, `isc_dsql_describe()`, and `isc_dsql_describe_bind()` statements can be used to determine the proper number of `XSQLVAR` structures to allocate, and the `XSQLDA_LENGTH` macro can be used to allocate the proper amount of space.

XSQLDA field descriptions

The following table describes the fields that comprise the `XSQLDA` structure:

Field definition	Description
<code>short version</code>	Indicates the version of the <code>XSQLDA</code> structure; set by an application The current version is defined in <code>ibase.h</code> as <code>SQLDA_VERSION1</code>
<code>char sqldaid[8]</code>	Reserved for future use
<code>ISC_LONG sqldabc</code>	Reserved for future use
<code>short sqln</code>	Indicates the number of elements in the <code>sqlvar</code> array; the application should set this field whenever it allocates storage for a descriptor
<code>short sqld</code>	Indicates the number of parameters for an input <code>XSQLDA</code> , or the number of select-list items for an output <code>XSQLDA</code> ; set by InterBase during an <code>isc_dsql_prepare</code> , <code>isc_dsql_describe()</code> , or <code>isc_dsql_describe_bind()</code> For an input descriptor, an <code>sqld</code> of 0 indicates that the SQL statement has no parameters; for an output descriptor, an <code>sqld</code> of 0 indicates that the SQL statement is not a SELECT statement
<code>XSQLVAR sqlvar</code>	The array of <code>XSQLVAR</code> structures; the number of elements in the array is specified in the <code>sqln</code> field

TABLE B.18 XSQLDA field descriptions

XSQLVAR field descriptions

The following table describes the fields that comprise the *XSQLVAR* structure:

Field definition	Description
<i>short sqltype</i>	Indicates the SQL datatype of parameters or select-list items; set by InterBase during <i>isc_dsql_prepare</i> , <i>isc_dsql_describe</i> (), or <i>isc_dsql_describe_bind</i> ()
<i>short sqlscale</i>	Provides scale, specified as a negative number, for exact numeric datatypes (DECIMAL, NUMERIC); set by InterBase during <i>isc_dsql_prepare</i> , <i>isc_dsql_describe</i> (), or <i>isc_dsql_describe_bind</i> ()
<i>short sqlsubtype</i>	Specifies the subtype for Blob data; set by InterBase during <i>isc_dsql_prepare</i> , <i>isc_dsql_describe</i> (), or <i>isc_dsql_describe_bind</i> ()
<i>short sqllen</i>	Indicates the maximum size, in bytes, of data in the <i>sqldata</i> field; set by InterBase during <i>isc_dsql_prepare</i> , <i>isc_dsql_describe</i> (), or <i>isc_dsql_describe_bind</i> ()
<i>char *sqldata</i>	For input descriptors, specifies either the address of a select-list item or a parameter; set by the application For output descriptors, contains a value for a select-list item; set by InterBase
<i>short *sqlind</i>	On input, specifies the address of an indicator variable; set by an application On output, specifies the address of column indicator value for a select-list item following a FETCH A value of 0 indicates that the column is not NULL; a value of -1 indicates the column is NULL; set by InterBase
<i>short sqlname_length</i>	Specifies the length, in bytes, of the data in field, <i>sqlname</i> ; set by InterBase during <i>isc_dsql_prepare</i> () or <i>isc_dsql_describe</i> ()
<i>char sqlname[32]</i>	Contains the name of the column. Not NULL (\0) terminated; set by InterBase during <i>isc_dsql_prepare</i> () or <i>isc_dsql_describe</i> ()
<i>short relname_length</i>	Specifies the length, in bytes, of the data in field, <i>relname</i> ; set by InterBase during <i>isc_dsql_prepare</i> () or <i>isc_dsql_describe</i> ()

TABLE B.19 XSQLVAR field descriptions

Field definition	Description
<i>char relname[32]</i>	Contains the name of the table; not NULL (\0) terminated, set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>
<i>short ownname_length</i>	Specifies the length, in bytes, of the data in field, <i>ownname</i> ; set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>
<i>char ownname[32]</i>	Contains the name of the table owner; not NULL (\0) terminated, set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>
<i>short aliasname_length</i>	Specifies the length, in bytes, of the data in field, <i>aliasname</i> ; set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>
<i>char aliasname[32]</i>	Contains the alias name of the column or the column name if no alias exists; not NULL (\0) terminated, set by InterBase during <i>isc_dsql_prepare()</i> or <i>isc_dsql_describe()</i>

TABLE B.19 XSQLVAR field descriptions (*continued*)

Index

A

- access mode parameter **65, 67**
- accessing
 - arrays **153–166**
 - arrays, DSQL applications **256**
- actions *See* events
- addresses
 - error messages **181, 182**
 - numeric data **95**
- aggregate functions
 - and arrays **153**
- ALIGN macro **95**
- allocating memory **90**
- animation **119**
- API calls, referencing/dereferencing **24, 25**
- API functions **23–30**
 - arrays **150**
 - BLOB data **118, 141**
 - categories, summarized **243–251**
 - conversions **167**
 - database **42**
 - DSQL **79**
 - error handling **29, 171**
 - DSQL applications **176–178**
 - events **187**
 - example programs **40**
 - informational **51, 131, 196**
 - processing SQL statements **83–85**
 - XSQLDAs and **85, 89**
 - prototypes **36**
 - transactions **60**
- applications
 - See also* DSQL applications
 - BLOB data and **119, 123, 128, 131**
 - allocating **288, 290**
 - type checking **137**
 - compiling *See* compiling
 - error handling **381, 382**
 - linking *See* linking
 - monitoring performance **298**
 - programming **23–30**
 - recovering data **364**
 - Windows *See* Windows applications
- arguments *See* parameters
- array buffers **158, 161**
- array columns **153, 163**
 - associating with arrays **165**
 - selecting **154, 156**
 - size, setting **263**
 - writing data to **160–166, 266**
- array descriptors **151–153**
 - creating **157, 160**
 - datatypes **261**
 - flags, setting **260, 263, 274**
 - initializing **273**
 - populating **152–153, 157, 260**
 - setting fields directly **152**
- array elements **150**
 - size, setting **260**
- array IDs **160, 255**
 - declaring **164**
 - fetching **158**
 - initializing **164**
- array slices **154–159**
 - changing one slice **161**
 - creating data buffers **161**
 - defined **150**
 - reading one slice **158**
 - writing data to **160–166**
- arrays **153–166**
 - See also* error status array
 - API functions **150**
 - associating with array columns **165**
 - column-major order, specifying **260, 263, 274**
 - creating **151, 164**
 - dropping **166**
 - DSQL applications and **83, 256**

- multi-dimensional **150, 151**
 - dimensions, range **150**
 - nested **150**
 - overview **150–151**
 - processing **28**
 - reading data **154–159**
 - creating data buffers **158**
 - retrieving data **151, 154, 156, 254**
 - selected rows **158**
 - row-major order, specifying **260, 263, 274**
 - subscripts **150**
 - subsets
 - retrieving **151**
 - writing to **151**
 - updating **160–166**
 - asynchronous events **189, 192**
 - canceling notification **197, 289**
 - requesting notification **369**
 - asynchronous traps **192–193**
 - creating **193**
 - attaching to a database **32**
 - attaching to databases **49–51, 80, 276**
 - See also* connecting to databases
 - DPBs and **44**
 - optional entries **24**
 - releasing system resources **305**
 - retrieving information about **51–57, 297**
 - temporarily **306**
 - Windows clients **32, 34**
 - automatic recovery functions **364**
- B**
- binary data **119**
 - binary file types, supported **119**
 - binary large objects *See* BLOB
 - bitmapped images **119**
 - BLOB (defined) **118**
 - BLOB API functions **118, 141**
 - BLOB columns
 - creating **128–130**
 - writing data to **123, 134**
 - BLOB filters and **142**
 - BLOB data **117, 120–131**
 - changing **126**
 - creating **288, 294**
 - deleting **130**
 - DSQL applications and **83, 120, 122–124, 127**
 - fetching **124**
 - filtering **142–147, 279**
 - processing **27, 287, 290**
 - reading **121–126, 134**
 - BLOB filters and **142**
 - retrieving information about **131–134, 282**
 - defaults **279**
 - selecting **121–124**
 - status messages **132**
 - storing **120, 128, 290**
 - support **119**
 - translating **136, 294**
 - updating **126, 127–128, 367**
 - BLOB datatype
 - arrays and **150**
 - NULL values **123, 126, 130**
 - user-defined **119**
 - BLOB descriptors **28, 134–135**
 - populating **135, 279**
 - structure, defined **134**
 - BLOB filter function
 - action macro definitions **141**
 - defining **137**
 - input/output fields **139**
 - BLOB filters **28, 134, 136–147**
 - control structure, defined **139**
 - external **136, 136–142**
 - declaring **137**
 - writing **137–142**
 - invoking **142**
 - NetWare servers and **136**
 - opening BLOBs **146, 288, 360**
 - specifying **294**
 - user-defined **136**
 - BLOB handles **129**
 - allocating **287**
 - BLOB IDs **120**
 - creating **129**
 - declaring **129**
 - resetting **126**
 - BLOB parameter buffers **142–146**
 - generating **281**

- numeric formats **145**
- parameter types **145**
- BLOB segments **135**
 - allocating **129**
 - defined **120**
 - reading **353**
 - retrieving **125**
 - size **284**
 - writing **367**
- BLOB subtypes **119**
 - filtering **360**
 - retrieving **279, 284**
 - setting **286**
- Borland C/C++ *See* C language
- BPBs *See* BLOB parameter buffers
- buffers **158, 161**
 - BLOB filters **142**
 - capturing error messages **174, 177**
 - database attachments **44, 51**
 - reinitializing event **196**
 - transactions **62**
- byte order, reversing **170**

C

- C language
 - converting dates **168–170**
 - directives **24, 27, 29**
 - error handling **180**
 - event notification **190**
 - predefined constants
 - item types (BLOBs) **131**
 - predefined macros **90**
 - action messages (BLOBs) **141**
 - datatypes (XSQLDAs) **90–95**
 - informational (SQL) **113**
 - time structures **168**
- CAD files **119**
- changes
 - undoing **375**
 - See also* rollbacks
- character sets
 - BLOB data and **135, 143**
 - retrieving current **279, 284**
 - setting **286**
- character strings

- converting
 - SQL statements to **84, 96**
- CHARACTER VARYING datatype
 - DSQL applications and **93**
- client applications *See* SQL client applications;
Windows clients
- client libraries **34**
- CLOSE **84**
- closing
 - databases **74**
- coercing datatypes (XSQLDAs) **94**
- column names
 - storing **152**
- column-major order
 - specifying **260, 263, 274**
- columns
 - arrays and **150, 154**
 - DSQL applications **312**
 - retrieving information about **312**
- commits **74, 76, 292**
 - See also* transactions
 - delaying **77**
 - executing **292, 362, 363**
 - retaining context **291**
- compiling **30**
- connecting to databases **42–51**
 - See also* attaching to databases
- constants
 - item types (BLOBs) **131**
- converting
 - dates **168–170**
- CREATE DATABASE **82**
 - DSQL applications **326**
- cursors
 - closing **334**
 - declaring **104, 111**
 - DSQL applications **334, 339**
 - naming **339**
 - opening **339**

D

- data
 - binary **119**
 - corrupted **74**
 - fetching **124, 158**

- DSQL applications **330**
- losing **74**
- protecting *See* security
- reading **154–159**
- recovering **364**
- retrieving **151, 154, 156, 254**
 - selected rows **158**
- storing **150, 163, 165**
- variable-length
 - processing in XSQLDAs **93, 94**
- database API functions **42**
- database handles **42–43**
 - assigning at run time **80**
 - declaring **25, 43**
 - defined **24**
 - initializing **25, 43**
 - multiple databases **80**
- database pages
 - retrieving information about **52**
- database parameter buffers **25, 41**
 - allocating storage space **48**
 - creating **44–47, 351**
 - expanding **351**
 - numeric formats **170**
 - parameters **45, 46**
 - adding at run time **48**
- databases
 - See also* multiple databases
 - accessing **41, 80**
 - attaching to **32**
 - changes, undoing **375**
 - closing **74**
 - creating **82**
 - dropping **306**
 - performance statistics **54–55**
 - referencing **24**
 - retrieving information about **52, 55, 392**
 - version numbers **392**
 - temporary **306**
- datatype coercion (XSQLDAs) **94**
- datatype macro constants (XSQLDAs) **90–95**
- datatypes **34, 36**
 - arrays **150, 261**
 - indeterminate **119**
- DATE data **167**
 - converting **168–170**
- DECIMAL datatype
 - DSQL applications and **93**
- DECLARE CURSOR **84**
- DECLARE FILTER **137**
- declaring
 - array IDs **164**
 - BLOB descriptors **28**
 - BLOB filters **137**
 - BLOB handles **129**
 - BLOB IDs **129**
 - cursors **104, 111**
 - database handles **25, 43**
 - error status vectors **172**
 - extended descriptor areas **85**
 - transaction handles **25, 61**
- default directories
 - Windows clients **34**
- defaults, retrieving
 - BLOBs **279**
- DELETE
 - BLOB data **130**
 - DSQL applications **340**
- deleting *See* dropping
- DESCRIBE **84**
- descriptor areas (extended) *See* XSQLDAs
- descriptor fields *See* array descriptors; BLOB
- descriptors
 - detaching from databases **43, 57**
 - See also* disconnecting from databases
- directories
 - Windows clients **34**
- disconnecting from databases **57, 305**
 - See also* detaching from databases
- displaying *See* output
- DLLs **40**
- DOS applications **33**
- DOS environment variables **33–34**
- DPBs *See* database parameter buffers
- dropping
 - arrays **166**
 - databases **58, 306**
- DSQL
 - programming methods **96–115**
- DSQL API functions **79**

- DSQL applications 27, 79
 - arrays and 83, 153, 256
 - BLOB processing 83, 120, 122–124, 127
 - closing cursors 334
 - declaring cursors 104, 111
 - defining cursors 339
 - error handling 176–178
 - extended descriptor areas *See* XSQLDAs
 - fetching data 330
 - queries 101–113
 - retrieving information about 314
 - SELECT statements 312
 - SQL statements 83–85
- DSQL statements
 - executing 317, 321, 325, 328
 - repeatedly 336
 - handles
 - allocating 308, 310
 - freeing 334
 - input parameters 314
 - retrieving information about 342
- DSQL_close option 334
- DSQL_drop option 335
- dynamic link libraries *See* DLLs
- dynamic SQL *See* DSQL

E

- environment variables 33–34
- environments
 - PCs 32
 - retrieving information about 54
- EPBs *See* event parameter buffers
- error codes 172, 181
 - examining status vector 178–182
 - system 182
 - translating 176, 381
- error messages 172
 - See also* SQL error messages
 - addresses 181, 182
 - building 355
 - displaying 173, 174, 176, 366
 - DSQL applications 176–178
- error status array 29
- error status vectors 29
 - checking 172

- clusters 178
- declaring 172
- numeric descriptors 179–181
 - numeric values 182
- parsing 178–187
- error-handling API functions 29, 171
 - DSQL applications 176–178
- error-handling routines 74, 181
 - SQL 381, 382
- errors 178, 375
 - run-time 171
 - transactions 74
- event buffer, reinitializing 196
- event parameter buffers 188
 - allocating 347
 - comparing event counts 349
 - creating 190
- events 187
 - asynchronous 189, 192
 - canceling notification 197, 289
 - requesting notification 369
 - posting 349
 - processing 29, 192–193
 - retrieving 196
 - synchronous 189, 394
 - transaction control 189
 - waiting on 191, 193
- example programs 40
- EXECUTE 84
- EXECUTE IMMEDIATE 84
- EXECUTE PROCEDURE
 - DSQL applications 312, 321, 326
- extended SQL descriptor areas *See* XSQLDAs
- external BLOB filters *See* BLOB filters

F

- FETCH 84
- fetching data 124, 158
 - See also* retrieving data
 - DSQL applications 330
- FILE structure 42
- file types
 - BLOB data 119
- filtering BLOB data 142–147, 279
 - See also* BLOB filters

formatting

error messages **174**

function prototypes **36**

G

GDS.DLL **40**

H

handles *See* database handles; transaction handles

header files *See* `ibase.h`

host names, specifying **34**

I

`ibase.h` **24, 25**

BLOB descriptors **28**

errors **29**

XSQLDA structures **27**

include files *See* `ibase.h`

indeterminate datatypes **119**

information item macros (SQL) **113**

informational API functions **51, 131, 196**

initializing

array descriptors **273**

array IDs **164**

BLOB handles **129**

BLOB IDs **129**

database handles **25, 43**

transaction handles **25, 61**

input descriptors *See* XSQLDAs

input fields

BLOB filters **140**

input parameters

DSQL statements **314**

INSERT

arrays **162, 165**

BLOB data **120, 127**

ISC_ARRAY_BOUND structure **151**

ISC_ARRAY_DESC structure **151**

`isc_array_get_slice()` **254**

`isc_array_lookup_bounds()` **152, 157, 161, 259**

`isc_array_lookup_desc()` **152, 263**

`isc_array_put_slice()` **160, 164, 266**

`isc_array_set_desc()` **152, 273**

`isc_attach_database()` **49, 276**

DPBs and **44**

`isc_blob_default_desc()` **279**

`isc_blob_gen_bpb()` **143, 281**

`isc_blob_info()` **131–134, 282**

`isc_blob_lookup_desc()` **284**

`isc_blob_set_desc()` **286**

`isc_cancel_blob()` **131, 287**

`isc_cancel_events()` **197, 289**

`isc_close_blob()` **290**

`isc_commit_retaining()` **75, 291**

`isc_commit_transaction()` **74, 76, 292**

`isc_create_blob2()` **129, 146, 294**

ISC_DATABASE environment variable **34**

`isc_database_info()` **51, 56, 297**

`isc_db_handle` type **43**

`isc_decode_date()` **169, 299, 300, 301**

`isc_detach_database()` **57, 305**

`isc_drop_database()` **58, 306**

`isc_dsql_allocate_statement()` **97, 308**

`isc_dsql_allocate_statement2()` **310**

`isc_dsql_describe()` **87, 312**

`isc_dsql_describe_bind()` **314**

`isc_dsql_exec_immed2()` **328**

`isc_dsql_execute()` **97, 317**

`isc_dsql_execute_immediate()` **96, 325**

`isc_dsql_execute2()` **321**

`isc_dsql_fetch()` **124, 158, 330**

SELECT statements and **331**

`isc_dsql_prepare()` **87, 97, 336**

`isc_dsql_sql_info()` and **342**

`isc_dsql_set_cursor_name()` **339**

`isc_dsql_sql_info()` **113–149, 342**

`isc_encode_date()` **170, 344, 345, 346**

`isc_event_block()` **190, 347**

`isc_event_counts()` **196–197, 349**

`isc_expand_dpb()` **48, 351**

ISC_EXPORT keyword **27**

`isc_get_segment()` **125, 353**

BLOB filters and **140**

`isc_info_truncated` value **131**

`isc_interprete()` **174, 355**

`isc_open_blob2()` **146, 360**

ISC_PASSWORD environment variable **34**

`isc_prepare_transaction()` **76, 362**

`isc_prepare_transaction2()` **77, 363**

isc_print_sqlerror() **176, 365**
 isc_print_status() **173, 366**
 isc_put_segment() **129, 367**
 BLOB filters and **139**
 isc_que_events() **192–195, 369**
 isc_rollback_transaction() **74, 77, 375**
 isc_sql_interprete() **177–178, 382**
 isc_sqlcode() **176, 381**
 isc_start_multiple() **71, 383**
 isc_start_transaction() vs. **387**
 isc_start_transaction() **69, 70, 386**
 ISC_STATUS pointer **173**
 isc_tr_handle type **61**
 isc_transaction_info() **389**
 ISC_USER environment variable **34**
 isc_vax_integer() **170, 391**
 isc_version() **392**
 isc_wait_for_event() **191–192, 394**
 isolation level parameter **65, 67**
 restrictive **66**
 item type constants (BLOBs) **131**
 item-list buffer **131**
 See also `isc_blob_info()`

L

libraries
 BLOB filtering routines **136**
 dynamic link *See* DLLs

limbo transactions **362**

linking **30**

lock resolution parameter **67**

log files **177**

M

macros
 action messages (BLOBs) **141**
 ALIGN **95**
 datatypes (XSQLDAs) **90–95**
 informational (SQL) **113**
 XSQLDA_LENGTH **90**

memory
 allocating **90**
 retrieving information about **54**

messages *See* error messages; status messages

Microsoft C/C++ *See* C language
 Microsoft Windows *See* Windows
 monitoring performance **298**
 multi-dimensional arrays **150, 151**
 dimensions, range **150**
 multiple databases
 attaching to **80**
 transactions and **71, 76, 362, 363, 383**

N

NCHAR VARYING datatype
 DSQL applications and **93**
 nested arrays **150**
 NetWare servers
 BLOB filters and **136**
 user names, returning **54**
 network DLLs **40**
 NULL pointers **69**
 NULL status **90**
 NULL values
 arrays and **153, 163, 164, 166**
 BLOB columns **123, 126, 130**
 BLOB handles **287**
 extended descriptor areas **92, 94**
 numbers **170**
 alignment addresses **95**
 byte ordering **170**
 processing in XSQLDAs **93, 94**
 NUMERIC datatype
 DSQL applications and **93**
 numeric values *See* values

O

ODS *See* on-disk structure

on-disk structure
 retrieving information about **53**

OPEN **84**

opening
 BLOBs **146, 288, 360**
 cursors **339**

output
 error messages **173, 174, 176, 365, 366**
 output descriptors *See* XSQLDAs
 output fields

- BLOB filters **140**
- P**
- parameters **351**
 - DPBs **45, 46, 48**
 - DSQL statements **314**
 - input **314**
 - SQL statements **85, 89, 90, 98, 106**
 - processing with no **96–97, 101–106**
 - transaction parameter buffers (TPBs) **59–72**
 - passwords **24**
 - See also* security
 - overriding **34**
 - supplying at run time **351**
 - Windows clients **32, 34**
 - PC development environments **32**
 - performance statistics **54–55**
 - performance, monitoring **298**
 - pointers
 - See also* cursors
 - FILE structure **42**
 - transactions **69**
 - PREPARE **84**
 - programming
 - API applications **23–30**
 - DSQL applications **79, 96–115**
 - error handling *See* errors
 - Windows applications **26**
 - protecting data *See* security
- Q**
- queries **27**
 - See also* SQL
 - arrays and **154**
 - DSQL applications **101–113**
- R**
- RDB\$TRANSACTIONS **293**
 - reading BLOB data **121–126, 134**
 - BLOB filters and **142**
 - reading data **154–159**
 - recovering data **364**
 - request buffer items **51**
 - request buffers
 - defined **51**
 - resetting BLOB IDs **126**
 - result buffers **51**
 - BLOBs **131–132**
 - defined **51**
 - retrieving data **124, 151, 154, 156, 254**
 - DSQL applications **330**
 - selected rows **158**
 - reversing byte order **170**
 - rollbacks **77, 375**
 - See also* transactions
 - routines
 - BLOB filters **136, 142**
 - error-handling **74, 181**
 - SQL **381, 382**
 - row-major order, specifying **260, 263, 274**
 - rows
 - DSQL applications **312**
 - retrieving information about **312**
 - run-time errors **171**
- S**
- security
 - attachment requirements **32**
 - Windows clients **34**
 - SELECT
 - See also* singleton SELECTs
 - BLOB data and **120, 121–124**
 - SELECT statements
 - arrays **153, 154–157**
 - DSQL applications **331**
 - executing **317, 321, 326**
 - preparing **337**
 - retrieving information about **312**
 - singleton SELECTs **329**
 - select-lists **101, 103, 107**
 - See also* queries
 - BLOB data **121**
 - defined **101**
 - processing items **105**
 - retrieving items **104**
 - SET TRANSACTION
 - DSQL applications **326**
 - signed numbers **170**
 - simultaneous transactions **66, 68**

- singleton SELECTs
 - DSQL applications **329**
- sound files, supported **119**
- SQL clients **32**
- SQL descriptor areas (extended) *See* XSQLDAs
- SQL error messages **176–178**
 - See also* SQLCODE variable
 - building **381, 382**
 - displaying **176, 365**
- SQL error-handling routines **381, 382**
- SQL statements
 - converting to character strings **84, 96**
 - creating **98, 102, 108**
 - DSQL applications and **83–85**
 - parameters, supplying values **89, 90, 98, 106**
 - executing **98, 101, 104, 106, 111, 113**
 - non-query statements and **96–101**
 - processing **98–101, 106–113**
 - with no parameters **96–97, 101–106**
 - retrieving select-list items **104**
 - selecting BLOB data **121–124**
- SQLCODE variable
 - DSQL applications **176**
 - return values **365**
- statements
 - retrieving **113–149, 342**
- status information **90, 172**
- status messages
 - BLOB data **132**
 - transactions **390**
- status vectors *See* error status vectors
- storing
 - BLOB data **120, 128**
 - data **150, 163, 165**
- string addresses
 - error messages **181, 182**
- strings *See* character strings
- subscripts (arrays) **150**
- sweeping databases
 - retrieving information about **54**
- synchronous events **189**
 - requesting notification **394**
- system crashes **364**
- system error codes **182**

T

- table names
 - storing **152**
- tables
 - accessing **68–69**
- temporary databases **306**
- text
 - BLOB type and **119**
- text files, supported **119**
- time structures **168**
- TPBs *See* transaction parameter buffers
- transaction handles **61–62**
 - assigning at run time **81**
 - declaring **25, 61**
 - defined **25**
 - initializing **25, 61**
- transaction IDs
 - tracking **389**
- transaction parameter buffers **26, 59**
 - constants **63**
 - creating **62–73**
 - default **69**
 - numeric formats **170**
- transactions **59**
 - access modes **65, 67**
 - accessing tables **68–69**
 - committing **74, 76, 77, 292**
 - executing two-phase commits **292, 362, 363**
 - retaining context **291**
 - ending **73**
 - events and **189**
 - isolation levels **65–67**
 - limbo **362**
 - locking conflicts **67**
 - multiple databases **71, 76, 362, 363, 383**
 - optimizing **75**
 - referencing **25**
 - retrieving information about **389**
 - rolling back **77, 375**
 - simultaneous **66, 68**
 - specifying attributes **62–73**
 - starting **25, 60, 383, 386**
 - status messages **390**
 - transaction existence blocks (TEBs) **71**

transaction parameter blocks (TPBs) 59–72
traps 192–193
See also events

U

unknown datatypes 119
unknown statements 113, 342
UPDATE
arrays 162–166
BLOB data 120, 127–128
DSQL applications 340
updating
arrays 160–166
BLOB data 126, 127–128, 367
user names 24
overriding 34
retrieving information about 54
supplying at run time 351
Windows clients 32, 34
user-defined BLOB filters 136
user-defined types
BLOB data 119

V

value parameters
SQL statements 85
values
See also NULL values
numeric descriptors 182
VARCHAR datatype
DSQL applications and 93, 94
variable-length data, processing 93, 94
vector-graphic files, supported 119
version numbers
databases 52, 392
on-disk structure 53
transaction processing 65
video files, supported 119
video segments 119
views
arrays and 153

W

Windows applications 26

defining datatypes 34, 36
event notification 191
setting default directories 34

Windows clients
attaching 32, 34
establishing program parameters 33
security 34
writing data
to arrays 160–166, 266
to BLOBs 123, 134
BLOB filters and 142

X

XSQLDA_LENGTH macro 90
XSQLDAs 85–95
See also XSQLVAR structure
address alignment 95
coercing datatypes 94
numbers 94
variable-length data 94
declaring 85
fields 87
input descriptors 85, 87, 89
allocating 90
arrays 162
creating 98, 106
output descriptors 85, 87, 89
allocating 90
creating 101, 107
resizing 337
retrieving NULL values 90, 92
select-list items and 101, 103, 107
arrays 154–155, 158
BLOB data 121
setting NULL values 94
specifying datatypes 90–95
numbers 93
variable-length data 93, 94
structures 27, 85
XSQLVAR structure 89, 100
allocating 87
arrays 156
BLOB data 122, 128
datatypes 92
defined 85

fields **88**
setting up select-list items **103**

