Operations Guide

VERSION 7

Borland®
InterBase®

Borland Software Corporation
100 Enterprise Way, Scotts Valley, CA 95066-3249
www.borland.com

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

COPYRIGHT © 2002 Borland Software Corporation. All rights reserved. All Borland brand and product names are trademarks or registered trademarks of Borland Software Corporation in the United States and other countries. Other product names are trademarks or registered trademarks of their respective holders.

Contents

Tables xi
Figures xiii

Chapter 1
Introduction
Who should use this guide 1-1
Topics covered in this guide 1-1
InterBase documentation 1-2
System requirements and server sizing 1-3
Primary InterBase features 1-3
SQL support 1-5
Multiuser database access 1-5
Transaction management 1-5
Multigenerational architecture 1-6
Optimistic row-level locking 1-6
Database administration 1-6
Managing server security 1-7
Backing up and restoring databases . . . 1-7
Maintaining a database 1-7
Viewing statistics 1-8
About InterBase SuperServer architecture . . . 1-8
How SuperServer differs from Classic architecture 1-8
Executable and processes 1-9
Lock management 1-9
Resource use 1-9
Local access method 1-9
Security on UNIX platforms 1-10
Overview of command-line tools 1-10
isql 1-10
gbak 1-10
gfix 1-11
gsec 1-11
gstat 1-11
iblockpr (gds_lock_print) 1-12
ibmgr 1-12

Chapter 2
IBConsole: The InterBase Interface
Starting IBConsole 2-2
IBConsole menus 2-2
Context menus 2-3
IBConsole toolbar 2-4
Tree pane 2-5
Work pane 2-6
Standard text display window 2-7
Switching between IBConsole windows . . . 2-7
Managing custom tools in IBConsole 2-8

Chapter 3
Server Configuration
Configuring server properties 3-1
The General tab 3-2
The Alias tab 3-3
SMP support 3-3
Expanded processor control: CPU_AFFINITY . 3-4
ibconfig parameter: MAX_THREADS 3-4
Using InterBase Manager to start and stop InterBase 3-5
Starting and stopping the InterBase Server on UNIX 3-6
Using ibmgr to start and stop the server . . . 3-6
Starting the server 3-6
Stopping the server 3-7
Starting the server automatically 3-7
The attachment governor 3-10
Using environment variables 3-11
ISC_USER and ISC_PASSWORD 3-11
The INTERBASE environment variables . . . 3-11
The TMP environment variable 3-12
Managing temporary files 3-12
Configuring history files 3-12
Configuring sort files 3-12
Configuring parameters in *ibconfig* 3-13
Viewing the server log file 3-17

Chapter 4
Network Configuration
Network protocols 4-1
Connecting to servers and databases 4-2
Registering a server 4-2
Logging in to a server 4-3
Logging out from a server 4-4
Unregistering a server 4-5
Registering a database 4-5
Connecting to a database 4-7
Connect 4-7
Connect as 4-7
Disconnecting a database 4-8
Unregistering a database 4-9
Connection-specific examples 4-9
Connection troubleshooting 4-10

Connection refused errors 4-10
Connection rejected errors. 4-12
Disabling automatic Internet dialup 4-13
Reorder network adapter bindings 4-13
Disabling autodial in the registry 4-14
Preventing RAS from dialing out for local network activity 4-14
Other errors 4-14
Communication diagnostics 4-15
DB Connection tab 4-16
To run a DB Connection test: 4-16
Sample output (local connection) 4-17
TCP/IP tab 4-17
NetBEUI tab. 4-19

Chapter 5
Database Security

Security model 5-1
The SYSDBA user 5-2
Other users 5-2
Users on UNIX 5-2
The InterBase security database 5-3
System table security 5-4
Older databases. 5-4
Scripts for changing database security 5-5
Migration issues 5-5
SQL privileges 5-6
Groups of users 5-6
SQL roles 5-6
UNIX groups 5-7
Other security measures 5-7
Restriction on using InterBase tools 5-7
Protecting your databases 5-8
User administration with IBConsole. 5-9
Displaying the User Information dialog . . . 5-9
Adding a user. 5-9
Modifying user configurations 5-10
Deleting a user 5-11
User administration with the InterBase API . .5-11
Using gsec to manage security 5-12
Running gsec remotely 5-12
Using gsec commands 5-12
Displaying the security database 5-13
Adding entries to the security database .5-13
Modifying the security database 5-14
Deleting entries from the security database 5-15
Using gsec from a Windows command prompt 5-15
gsec error messages. 5-16

Chapter 6
Database Configuration and Maintenance

Database files 6-1
Database file size 6-2
Dynamic file sizing 6-2
External files 6-2
Temporary files 6-3
File naming conventions 6-3
Primary file specifications 6-3
Secondary file specifications 6-4
Multifile databases 6-4
Adding database files 6-4
Altering database file sizes 6-4
Maximum number of files 6-5
Application considerations 6-5
Reorganizing file allocation. 6-5
Networked file systems. 6-5
On-disk structure (ODS) 6-6
Read-write and read-only databases 6-6
Read-write databases 6-6
Read-only databases 6-7
Properties of read-only databases 6-7
Making a database read-only. 6-7
Read-only with older InterBase versions . 6-8
Creating databases 6-8
Database options 6-10
Page size 6-10
Default character set. 6-10
SQL dialect 6-10
Dropping databases 6-10
Backup file properties 6-11
Removing database backup files 6-12
Shadowing. 6-12
Tasks for shadowing 6-12
Advantages of shadowing 6-13
Limitations of shadowing 6-13
Creating a shadow 6-14
Creating a single-file shadow. 6-14
Creating a multifile shadow 6-15
Auto mode and manual mode 6-16
Conditional shadows 6-17
Activating a shadow 6-17
Dropping a shadow 6-17
Adding a shadow file 6-18
Setting database properties 6-18
Alias tab. 6-18

General tab 6-19
Sweep interval and automated housekeeping .6-21
Overview of sweeping 6-21
Setting the sweep interval 6-22
Disabling automatic sweeping 6-23
Performing an immediate database sweep .6-23
Configuring the database cache 6-24
Default cache size per database. 6-24
Default cache size per ISQL connection . . .6-25
Setting cache size in applications 6-25
Default cache size per server 6-25
Verifying cache size 6-25
Forced writes vs. buffered writes 6-26
Validation and repair 6-26
Validating a database 6-27
Validating a database using gfix 6-27
Validating a database using IBConsole . .6-27
Repairing a corrupt database 6-29
Shutting down and restarting databases 6-30
Shutting down a database 6-31
Shutdown timeout options 6-31
Shutdown options 6-31
Restarting a database 6-33
Limbo transactions. 6-33
Recovering transactions 6-34
Transaction tab 6-34
Details tab 6-35
Viewing the administration log 6-36
gfix command-line tool 6-37
gfix error messages 6-39

Chapter 7
Licensing

Software activation certificates 7-1
License registration tools 7-1
Using IBConsole 7-2
Viewing existing licenses 7-2
Adding a certificate ID/key 7-2
Deleting a certificate ID/key 7-3
Command-line registration utility 7-3
Available certificates 7-5
A note on simultaneous connections. 7-5
The InterBase license file 7-6
Options in the *ib_license.dat* file 7-6

Chapter 8
Database Backup and Restore

Benefits of backup and restore 8-1
Database ownership 8-2
Backing up a database using IBConsole 8-2
Backup options 8-4
Format 8-5
Metadata Only 8-5
Garbage collection 8-5
Transactions in limbo 8-6
Checksums 8-6
Convert to Tables 8-6
Verbose Output 8-7
Transferring databases to servers running other operating systems. 8-7
Restoring a database using IBConsole 8-8
Restore options 8-10
Page Size 8-10
Overwrite 8-11
Commit After Each Table 8-11
Create Shadow Files 8-11
Deactivate Indexes. 8-12
Validity Conditions 8-12
Use All Space 8-12
Verbose Output 8-13
gbak command-line tool 8-13
Database backup 8-13
Backing up a database with gbak 8-15
Restoring a database with gbak 8-16
Using gbak with InterBase Service Manager 8-19
The user name and password 8-20
Some backup and restore examples 8-20
Backup examples 8-20
gbak error messages 8-22

Chapter 9
Database Statistics and Connection Monitoring

Monitoring with system temporary tables 9-1
Querying system temporary tables 9-2
Refreshing the temporary tables 9-2
Listing the temporary tables 9-3
Security 9-3
Examples 9-3
Updating system temporary tables 9-4
Making global changes 9-4
Viewing statistics using IBConsole 9-5
Database statistics options 9-6
All Options 9-6
Data Pages 9-6
Database Log 9-7
Header Pages. 9-8
Index Pages. 9-10

System Relations 9-11
Monitoring client connections with IBConsole . 9-11
The gstat command-line tool 9-12
Viewing lock statistics 9-15
Retrieving statistics with *api_database_info*() . . 9-17

Chapter 10
Interactive Query

IBConsole ISQL 10-1
The interactive SQL window 10-1
ISQL menus 10-2
ISQL toolbar 10-6
ISQL work areas 10-7
Status bar 10-7
Managing ISQL temporary files 10-8
Executing SQL statements 10-8
Executing SQL interactively 10-8
Preparing SQL statements 10-9
Valid SQL statements 10-9
Loading and executing a SQL script file . 10-9
Committing and rolling back transactions 10-10
Saving ISQL input and output 10-10
Saving SQL input 10-10
Saving SQL output 10-10
Changing ISQL settings 10-11
Options tab 10-11
Advanced tab 10-13
Inspecting database objects 10-14
Viewing object properties 10-15
Viewing metadata 10-15
Extracting metadata 10-16
Extracting metadata 10-18
Command-line isql tool 10-19
Invoking isql 10-19
Command-line options 10-20
Using warnings 10-21
Examples 10-22
Exiting isql 10-22
Connecting to a database 10-22
Setting isql client dialect 10-23
Transaction behavior in isql 10-24
Extracting metadata 10-24
isql Commands 10-25
SHOW commands 10-26
SET commands 10-26
Other isql commands 10-26
Exiting isql 10-26
Error handling 10-26
isql command reference 10-27
BLOBDUMP 10-28
EDIT 10-28
EXIT 10-29
HELP 10-29
INPUT 10-30
OUTPUT 10-31
QUIT 10-32
SET 10-32
SET AUTODDL 10-34
SET BLOBDISPLAY 10-35
SET COUNT 10-37
SET ECHO 10-38
SET LIST 10-39
SET NAMES 10-40
SET PLAN 10-40
SET STATS 10-41
SET TERM 10-42
SET TIME 10-43
SHELL 10-44
SHOW CHECK 10-45
SHOW DATABASE 10-45
SHOW DOMAINS 10-46
SHOW EXCEPTIONS 10-47
SHOW FILTERS 10-47
SHOW FUNCTIONS 10-48
SHOW GENERATORS 10-49
SHOW GRANT 10-49
SHOW INDEX 10-50
SHOW PROCEDURES 10-51
SHOW ROLES 10-52
SHOW SYSTEM 10-52
SHOW TABLES 10-53
SHOW TRIGGERS 10-53
SHOW VERSION 10-54
SHOW VIEWS 10-55
Using SQL scripts 10-55
Creating an isql script 10-55
Running a SQL script 10-56
Committing work in a SQL script 10-56
Adding comments in an isql script 10-57

Chapter 11
Database and Server Performance

Introduction 11-1
Hardware configuration 11-2
Choosing a processor speed 11-2
Sizing memory 11-2
Using high-performance I/O subsystems . 11-3
Distributing I/O 11-4

Using RAID 11-4
Using multiple disks for database files . . 11-4
Using multiple disk controllers 11-5
Making drives specialized 11-5
Using high-bandwidth network systems . . 11-5
Using high-performance bus 11-6
Useful links 11-7
Operating system configuration 11-7
Disabling screen savers 11-8
Console logins 11-8
Sizing a temporary directory 11-9
Use a dedicated server 11-9
Optimizing Windows NT for network applications 11-9
Understanding Windows NT pitfalls . . . 11-10
Network configuration 11-10
Choosing a network protocol 11-11
NetBEUI 11-11
TCP/IP 11-11
Configuring hostname lookups 11-11
Database properties 11-12
Choosing a database page size 11-12
Setting the database page fill ratio 11-13
Sizing database cache buffers 11-14
Buffering database writes 11-15
Database design principles 11-15
Defining indexes 11-16
What is an index? 11-16
What queries use an index? 11-16
What queries don't use indexes? 11-17
Directional indexes 11-17
Normalizing databases 11-17
Choosing Blob segment size 11-17
Database tuning tasks 11-18
Tuning indexes 11-18
Rebuilding indexes 11-18
Recalculating index selectivity 11-18
Performing regular backups 11-18
Increasing backup performance 11-18
Increasing restore performance 11-19
Facilitating garbage collection 11-19
Application design techniques 11-19
Using transaction isolation modes 11-19
Using correlated subqueries 11-20
Preparing parameterized queries 11-21
Designing query optimization plans 11-22
Deferring index updates 11-22
Application development tools 11-22
InterBase Express™ (IBX) 11-23
IB Objects 11-23
Borland Database Engine 11-23
BDE driver flags 11-23
SQL passthru mode 11-24
SQL query mode 11-24
Visual components 11-25
Understanding fetch-all operations . . . 11-25
TQuery 11-25
TTable 11-26
InterBase Limits 11-26

Chapter 12
Data Replication

About IBReplicator 12-1
Configuration tools 12-2
Replication server 12-2
InterBase support 12-2
Requirements 12-2
Windows system requirements 12-2
Sun Solaris 2.5.x or 2.6.x system requirements 12-3
Linux system requirements 12-3
IBReplicator features 12-3
Overview 12-3
InterBase advanced features 12-3
Ease of use 12-4
Components of IBReplicator 12-4
About data replication 12-5
How IBReplicator works 12-6
The basic sequence 12-7
IBReplicator's strategies 12-7
Resolving precedence issues 12-8
Replication involving newer InterBase datatypes 12-8
Operation logging 12-9
Viewing schema 12-9
Running the Replication Server 12-9
Windows platforms 12-9
UNIX and Linux platforms 12-10
Using IBReplicator 12-11
Using Replication Manager 12-11
Managing configuration databases 12-12
Creating a configuration database . . . 12-12
Working with configurations 12-13
The default configuration 12-13
Registering databases 12-14
Defining replication schemata 12-15
Defining a replication 12-15
Customizing default settings 12-16

Creating schemata 12-17
To define a new schema 12-17
To edit a schema 12-18
To delete a schema 12-18
The schema node 12-18
Choosing the source database. 12-18
Choosing the target databases 12-20
Specifying a target database 12-20
Cloning targets 12-21
Choosing replicated tables and stored procedures. 12-21
Specifying replicated tables 12-21
Removing replicated tables 12-22
Table settings and row-level replication. 12-22
Using WHERE conditions to identify rows 12-22
Replicating to stored procedures 12-23
Choosing replicated columns 12-26
Defining primary keys 12-26
Identifying data columns 12-26
Creating system objects 12-27
Special user REPL 12-27
Bi-Directional replication 12-27
n-Way replication. 12-28
Problems and workarounds 12-29

Appendix A
Migrating to InterBase 6 and later

Migration process A-1
Server and database migration A-2
Client migration A-2
Migration Issues A-2
InterBase SQL dialects A-2
Clients and databases A-3
Keywords used as identifiers A-3
Understanding SQL dialects A-3
Dialect 1 clients and databases A-3
Dialect 2 clients A-4
Dialect 3 clients and databases A-4
Setting SQL dialects A-5
Setting the isql client dialect. A-5
Setting the gpre dialect A-6
Setting the database dialect A-6
Features and dialects A-7
Features available in all dialects A-7
IBConsole, InterBase's graphical interface A-7
Read-only databases A-7
Altering column definitions A-7
Altering domain definitions A-7
The EXTRACT() function. A-7
SQL warnings A-7
The Services API, Install API, and Licensing API. A-8
New gbak functionality A-8
InterBase Express™ (IBX). A-8
Features available only in dialect 3 databases A-8
Delimited identifiers. A-8
INT64 data storage. A-8
DATE and TIME datatypes. A-8
New InterBase keywords. A-9
Delimited identifiers A-9
How double quotes have changed. . . . A-10
DATE, TIME, and TIMESTAMP datatypes . . A-10
Converting TIMESTAMP columns to DATE or TIME A-12
Casting date/time datatypes A-12
Adding and subtracting datetime datatypes A-14
Using date/time datatypes with aggregate functions A-16
Default clauses A-16
Extracting date and time information . . A-16
DECIMAL and NUMERIC datatypes A-18
Compiled objects A-19
Generators A-20
Miscellaneous issues A-20
Migrating servers and databases A-20
"In-place" server migration A-21
Migrating to a new server A-22
About InterBase 6, dialect 1 databases . . . A-23
Migrating databases to dialect 3. . . . . A-24
Overview A-24
Method one: in-place migration A-24
Column defaults and column constraints A-27
Unnamed table constraints A-28
About NUMERIC and DECIMAL datatypes A-28
Method two: migrating to a new database . A-30
Migrating older databases A-31
Migrating clients A-31
IBReplicator migration issues A-33
Migrating data from other DBMS products . . A-33

Appendix B
InterBase Limits
Various InterBase limits B-2

Index I-1

Tables

1.1 InterBase features 1-3
1.2 Comparison of Classic and SuperServer architectures 1-8
2.1 IBConsole context menu for a server icon 2-3
2.2 IBConsole context menu for a connected database icon 2-4
2.3 IBConsole standard toolbar 2-4
2.4 Server/database tree commands 2-6
3.1 Contents of *ibconfig* 3-14
4.1 Matrix of connection supported protocols 4-1
4.2 Using Communication Diagnostics to diagnose connection problems 4-18
5.1 Format of the InterBase security database 5-4
5.2 Summary of gsec commands 5-13
5.3 gsec options 5-14
5.4 gsec security error messages. 5-16
6.1 Auto vs. manual shadows 6-16
6.2 General options 6-20
6.3 Validation options 6-29
6.4 gfix options 6-37
6.5 gfix database maintenance error messages 6-39
7.1 iblicense commands and their options. . 7-4
7.2 iblicense options 7-4
7.3 Certificate keys available for InterBase. . 7-5
7.4 InterBase components 7-6
7.5 InterBase license options. 7-7
8.1 gbak arguments 8-14
8.2 gbak backup options 8-14
8.1 gbak restore options 8-17
8.2 host_service syntax for calling the Service Manager with gbak 8-19
8.3 gbak backup and restore error messages. 8-22
9.1 InterBase temporary system tables 9-2
9.2 Data page information 9-7
9.3 Header page information 9-8
9.4 Index pages information 9-10
9.5 gstat options 9-13
9.6 iblockpr/gds_lock_print options 9-16
9.7 Database I/O statistics information items . . 9-17
10.1 Toolbar buttons for executing SQL statements 10-6
10.2 Options Tab of the SQL Options dialog 10-12
10.3 Advanced tab of the SQL Options dialog 10-13
10.4 Object inspector toolbar buttons 10-15
10.5 Metadata information items 10-16
10.6 Metadata extraction constraints 10-17
10.7 Order of metadata extraction. 10-18
10.8 isql command-line options 10-20
10.9 isql extracting metadata arguments . . 10-24
10.10 SQLCODE and message summary . . 10-27
10.11 isql commands 10-27
10.12 SET statements 10-33
11.1 Matrix of BDE driver flags values 11-24
11.2 InterBase specifications 11-27
12.1 IBReplicator configuration tools 12-2
A.1 isql dialect precedence A-5
A.2 Results of casting to date/time datatypes A-13
A.3 Results of casting to date/time datatypes A-14
A.4 Adding and subtracting date/time datatypes A-15
A.5 Extracting date and time information . . A-17
A.6 Handling quotation marks inside of strings A-25
A.7 Migrating clients: summary A-32
B.1 InterBase specifications B-2

Figures

2.1 IBConsole window 2-2
2.2 IBConsole Toolbar 2-4
2.3 IBConsole Tree pane 2-5
2.4 Active Windows dialog 2-7
2.5 Tools dialog 2-8
2.6 Tool Properties dialog 2-8
3.1 Server Log dialog 3-18
4.1 Register Server and Connect dialog 4-2
4.2 Server Login dialog 4-4
4.3 Register Database and Connect dialog . . 4-6
4.4 Database connect dialog 4-8
4.5 Communications dialog - DB Connection 4-16
4.6 Communications dialog - TCP/IP 4-17
4.7 Communications dialog - NetBEUI 4-19
5.1 User information dialog 5-9
6.1 Create Database dialog 6-9
6.2 Backup alias properties 6-11
6.3 Database Properties - Alias tab 6-19
6.4 Database Properties - General tab 6-20
6.5 Database Validation dialog 6-28
6.6 Validation report dialog 6-29
6.7 Database shutdown dialog 6-31
6.8 Transaction Recovery - limbo transactions 6-35
6.9 Transaction recovery: Details 6-36
6.10 Administration Log dialog 6-36
7.1 The Add Certificate dialog 7-2
8.1 Database backup dialog 8-3
8.2 Database backup options 8-5
8.3 Database backup verbose output 8-7
8.4 Database Restore dialog 8-8
8.5 Database restore options 8-10
8.6 Database restore verbose output 8-13
9.1 Database Statistics options 9-5
9.2 Database Statistics dialog 9-6
9.3 Database connections dialog 9-12
10.1 The interactive SQL editor in IBConsole . 10-2
10.2 Options tab of the SQL Options dialog . . 10-11
10.3 Advanced tab of the SQL Options dialog 10-13
10.4 Object inspector 10-14
11.1 Comparing external transfer rate of disk I/O interfaces 11-3
11.2 Comparing bandwidth of network interfaces 11-6
11.3 Comparing throughput of bus technologies 11-7
12.1 The Configuration Editor 12-13
12.2 Replication Manager database tab . . . 12-14
12.3 Replication Manager Replications tab . 12-16

Chapter

1

Introduction

The InterBase *Operations Guide* is a task-oriented reference of procedures to install, configure, and maintain an InterBase database server or Local InterBase workstation.

This chapter describes who should read this book, and provides a brief overview of the capabilities and tools available in the InterBase product line.

Who should use this guide

The InterBase *Operations Guide* is for database administrators or system administrators who are responsible for operating and maintaining InterBase database servers. The material is also useful for application developers who wish to understand more about InterBase technology. The guide assumes knowledge of:

- Server operating systems including Windows NT/2000, Linux, and UNIX
- Networks and network protocols
- Application programming

Topics covered in this guide

- Introduction to InterBase features
- Server configuration, startup and shutdown
- Network configuration and troubleshooting guidelines

- Security configuration for InterBase servers, databases, and data; reference for the security configuration tools
- Database configuration and maintenance options; reference for the maintenance tools
- Backing up and restoring databases; reference for the backup tools
- Performance troubleshooting and tuning guidelines.
- Database and server statistics monitoring
- Interactive query profiling; reference for the interactive query tools

InterBase documentation

The complete documentation set for InterBase 7.0 is included with the product in PDF format. It does not install by default. When you run the install you must choose "Custom" explicitly choose to install them. They require about 13MB of space on your disk. You can also install them at any later time, access them from the CD-ROM, or copy them from the CD-ROM.

The books are available for purchase in printed form from **http://shop.borland.com**.

The document set has been indexed for use with Acrobat's Full Text Search. You can view it with either Acrobat Reader 4 or 5. If you have Acrobat Reader With Search (included on the InterBase CD-ROM), your interface will have one button for Find and another for Search. The Find button looks only within the current document. The Search button looks though the entire document set.

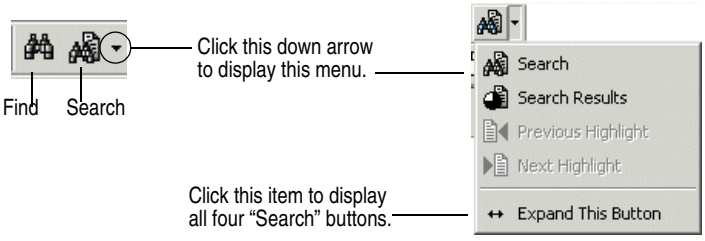

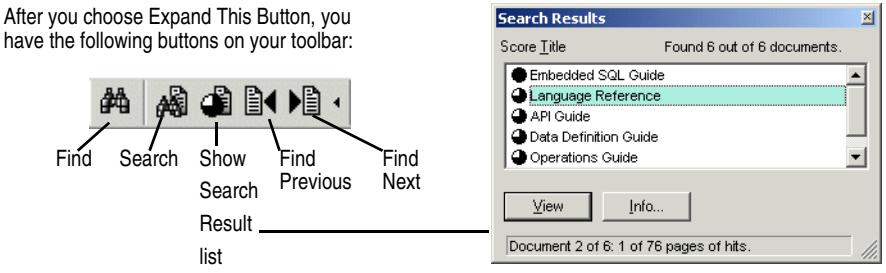

System requirements and server sizing

InterBase server runs on a variety of platforms, including Microsoft Windows NT, Windows 2000, and Windows 98/ME, Linux, and several UNIX operating systems.

The InterBase server software makes efficient use of system resources on the server node. The server process uses little more than 1.9MB of memory. Typically, each client connection to the server adds approximately 115KB of memory. This varies based on the nature of the client applications and the database design, so the figure is only a baseline for comparison.

The minimal software installation requires disk space ranging from 9MB to 12MB, depending on platform. During operation, InterBase's sorting routine requires additional disk space as scratch space. The amount of space depends on the volume and type of data the server is requested to sort.

The InterBase client also runs on any of these operating systems. In addition, InterBase provides the InterClient Java client interface using the JDBC standard for database connectivity. Client applications written in Java can run on any client platform that supports Java, even if InterBase does not explicitly list it among its supported platforms. Examples include the Macintosh and Internet appliances with embedded Java capabilities.

Primary InterBase features

InterBase offers all the benefits of a full-featured RDBMS. The following table lists some of the key InterBase features:

Table 1.1 InterBase features

Feature	Description
Network protocol support	• All platforms of InterBase support TCP/IP • InterBase servers and clients for Windows support NetBEUI/named pipes
SQL-92 entry-level conformance	ANSI standard SQL, available through an Interactive SQL tool and Borland desktop applications
Simultaneous access to multiple databases	One application can access many databases at the same time
multigenerational architecture	Server maintains older versions of records (as needed) so that transactions can see a consistent view of data
Optimistic row-level locking	Server locks only the individual records that a client updates, instead of locking an entire database page

Table 1.1 InterBase features (*continued*)

Feature	Description
Query optimization	Server optimizes queries automatically, or you can manually specify a query plan
Blob datatype and Blob filters	Dynamically sizeable datatypes that can contain unformatted data such as graphics and text
Declarative referential integrity	Automatic enforcement of cross-table relationships (between FOREIGN and PRIMARY KEYs)
Stored procedures	Programmatic elements in the database for advanced queries and data manipulation actions
Triggers	Self-contained program modules that are activated when data in a specific table is inserted, updated, or deleted
Event alerters	Messages passed from the database to an application; enables applications to receive asynchronous notification of database changes
Updatable views	Views can reflect data changes as they occur
User-defined functions (UDFs)	Program modules that run on the server
Outer joins	Relational construct between two tables that enables complex operations
Explicit transaction management	Full control of transaction start, commit, and rollback, including named transactions
Concurrent multiple application access to data	One client reading a table does not block others from it
multidimensional arrays	Column datatypes arranged in an indexed list of elements
Automatic two-phase commit	Multi-database transactions check that changes to all databases happen before committing (InterBase Server only)
InterBase API	Functions that enable applications to construct SQL/DSQL statements directly to the InterBase engine and receive results back
gpre	Preprocessor for converting embedded SQL/DSQL statements and variables into a format that can be read by a host-language compiler; included with the InterBase server license
IBConsole	Windows tool for data definition, query, database backup, restoration, maintenance, and security
isql	Command-line version of the InterBase interactive SQL tool; can be used instead of IBConsole.

Table 1.1 InterBase features (*continued*)

Feature	Description
Command-line database administrator utilities	Command-line version of the InterBase database administration tools; can be used instead of IBConsole
Header files	Files included at the beginning of application programs that define InterBase datatypes and function calls
Example make files	Files that demonstrate how to invoke the makefiles to compile and link InterBase applications
Example programs	C programs, ready to compile and link, which you can use to query standard InterBase example databases on the server
Message file	*interbase.msg*, containing messages presented to the user

SQL support

InterBase conforms to entry-level SQL-92 requirements. It supports declarative referential integrity with cascading operations, updatable views, and outer joins. InterBase Server provides libraries that support development of embedded SQL and DSQL client applications. On all InterBase platforms, client applications can be written to the InterBase API, a library of functions with which to send requests for database operations to the server.

InterBase also supports extended SQL features, some of which anticipate SQL99 extensions to the SQL standard. These include stored procedures, triggers, SQL roles, and segmented Blob support.

For information on SQL, see the *Language Reference*.

Multiuser database access

InterBase enables many client applications to access a single database simultaneously. A client applications can also access the multiple databases simultaneously. SQL triggers can notify client applications when specific database events occur, such as insertions or deletions.

You can write user-defined functions (UDFs) and store them in an InterBase database, where they are accessible to all client applications accessing the database.

Transaction management

Client applications can start multiple simultaneous transactions. InterBase provides full and explicit transaction control for starting, committing, and rolling back transactions. The statements and functions that control starting a transaction also control transaction behavior.

InterBase transactions can be isolated from changes made by other concurrent transactions. For the life of these transactions, the database appears to be unchanged except for the changes made by the transaction. Records deleted by another transaction exist, newly stored records do not appear to exist, and updated records remain in the original state.

For information on transaction management, see the *Embedded SQL Guide*.

Multigenerational architecture

InterBase provides expedient handling of time-critical transactions through support of data concurrency and consistency in mixed use—query and update—environments. InterBase uses a multigenerational architecture, which creates and stores multiple versions of each data record. By creating a new version of a record, InterBase allows all clients to read a version of any record at any time, even if another user is changing that record. InterBase also uses transactions to isolate groups of database changes from other changes.

Optimistic row-level locking

Optimistic locks are applied only when a client actually updates data, instead of at the beginning of a transaction. InterBase uses optimistic locking technology to provide greater throughput of database operations for clients.

InterBase implements true row-level locks, to restrict changes only to the records of the database that a client changes; this is distinct from page-level locks, which restrict any arbitrary data that is stored physically nearby in the database. Row-level locks permit multiple clients to update data that is in the same table without coming into conflict. This results in greater throughput and less serialization of database operations.

InterBase also provides options for pessimistic table-level locking. See the *Embedded SQL Guide* for details.

Database administration

InterBase provides both GUI and command-line tools for managing databases and servers. You can perform database administration on databases residing on Local InterBase or InterBase Server with IBConsole, a Windows application running on a client PC. You can also use command-line database administration utilities on the server.

IBConsole and command-line tools enable the database administrator to:

- Manage server security
- Back up and restore a database
- Perform database maintenance

- View database and lock manager statistics

You can find more information on server security later in this chapter, and later chapters describe individual tasks you can accomplish with IBConsole and the command-line tools.

Managing server security

InterBase maintains a list of user names and passwords in a security database. The security database allows clients to connect to an InterBase database on a server if a user name and password supplied by the client match a valid user name and password combination in the InterBase security database (*admin.ib* by default), on the server.

You can add and delete user names and modify a user's parameters, such as password and user ID.

For information about managing server security, see **Chapter 5, "Database Security."**

Backing up and restoring databases

You can backup and restore a database using IBConsole or command-line **gbak**. A backup can run concurrently with other processes accessing the database because it does not require exclusive access to the database.

Database backup and restoration can also be used for:

- Erasing obsolete versions of database records
- Changing the database page size
- Changing the database from single-file to multifile
- Transferring a database from one operating system to another
- Backing up only a database's metadata to recreate an empty database

For information about database backup and recovery, see **Chapter 8, "Database Backup and Restore."**

Maintaining a database

You can prepare a database for shutdown and perform database maintenance using either IBConsole or the command-line utilities. If a database incurs minor problems, such as an operating system write error, these tools enable you to sweep a database without taking the database off-line.

Some of the tasks that are part of database maintenance are:

- Sweeping a database
- Shutting down the database to provide exclusive access to it
- Validating table fragments
- Preparing a corrupt database for backup

- Resolving transactions “in limbo” from a two-phase commit
- Validating and repairing the database structure

For information about database maintenance, see Chapter 6, “Database Configuration and Maintenance.”

Viewing statistics

You can monitor the status of a database by viewing statistics from the database header page, and an analysis of tables and indexes. For more information, see Chapter 9, “Database Statistics and Connection Monitoring.”

About InterBase SuperServer architecture

InterBase uses SuperServer architecture: a multi-client, multi-threaded implementation of the InterBase server process. SuperServer replaces the Classic implementation used for previous versions of InterBase. SuperServer serves many clients at the same time, using threads instead of separate server processes for each client. Multiple threads share access to a single server process.

How SuperServer differs from Classic architecture

If you are upgrading from a previous version of InterBase, it will help to understand how the differences listed in the table below affect database operations. The differences are explained in the paragraphs that follow the table.

Table 1.2 Comparison of Classic and SuperServer architectures

Feature	Classic	SuperServer
Executable	**gds_inet_server**	ibserver
Processes	Multiple, on demand, one instance per client	single, threaded, one thread per client
Lock management	**gds_lock_mgr** process	Implemented as a thread
Resource use	One cache per process	One cache space for all clients
Local access	Direct I/O	I/O by proxy (network style)
Security on UNIX	Executable and Lock Manager	Can run as non-root **uid**, with access restrictions

Executable and processes

- **Classic**: Runs on demand as multiple processes. When a client attempts to connect to an InterBase database, one instance of the **gds_inet_server** executable runs and remains dedicated to that client connection for the duration of the connection.
- **SuperServer**: Runs as a single process, an invocation of the **ibserver** executable. **ibserver** is started once by the system administrator or by a system boot script. This process runs always, waiting for connection requests. Even when no client is connected to a database on the server, **ibserver** continues to run quietly. The SuperServer process does not depend on **inetd**; it waits for connection requests to the **gds_db** service itself.

Lock management

- **Classic**: For every client connection a separate server process is started to execute the database engine, and each server process has a dedicated database cache. The server processes contend for access to the database, so a Lock Manager subsystem is required to arbitrate and synchronize concurrent page access among the processes.
- **SuperServer**: The lock manager is implemented as a thread in the **ibserver** executable. Therefore, InterBase does not use the **gds_lock_mgr** process. Likewise, POSIX signals are not used by the lock manager thread in SuperServer; interthread communication mechanisms are used instead.

Resource use

- **Classic**: Each instance of **gds_inet_server** keeps a cache of database pages in its memory space. While the resource use per client is greater than in SuperServer, Classic uses less overall resources when the number of concurrent connections is low.
- **SuperServer**: One cache space is used for all client attachments, allowing more efficient use of cache memory.

Local access method

- **Classic**: Permits application processes that are running on the same host as the database to perform I/O on database files directly.
- **SuperServer**: Requires applications to request **ibserver** I/O operations by proxy, using a network method. SuperServer currently does not support true local access as Classic does, even when the client application and **ibserver** are running on the same host. This is important, because Classic's use of direct I/O is faster than the interprocess communication method that is required when using SuperServer.

Security on UNIX platforms

- **Classic**: Server executables **gds_inet_server** and **gds_lock_mgr** must run as root. The lock manager must have the superuser privilege to send signals to the processes.
- **SuperServer**: Can be configured to run as a non-root **uid**, for enhanced security. The permissions on database files can be restricted to allow only InterBase server **uid** to access the database.

Overview of command-line tools

For each task that you can perform in IBConsole, there is a command-line tool that you can run in a command window or console to perform the same task.

The UNIX versions of InterBase include all of the following command-line tools. The graphical Windows tools do not run on a UNIX workstation, though you can run most of the tools on Windows to connect to and operate on InterBase databases that reside on UNIX servers.

An advantage of noninteractive, command-line tools is that you can use them in batch files or scripts to perform common database operations. You can automate execution of scripts through your operating system's scheduling facility (**cron** on UNIX, **AT** on Windows NT/2000). It is more difficult to automate execution of graphical tools.

isql

The **isql** tool is a shell-type interactive program that enables you to quickly and easily enter SQL statements to execute with respect to a database. This tool uses InterBase's Dynamic SQL mechanism to submit a statement to the server, prepare it, execute it, and retrieve any data from statements with output (for example, from a SELECT or EXECUTE PROCEDURE). **isql** manages transactions, displays metadata information, and can produce and execute scripts containing SQL statements.

See Chapter 10, "Interactive Query" for full documentation and reference on **isql** and using **isql** from IBConsole.

gbak

The **gbak** tool provides options for backing up and restoring databases. **gbak** now backs up to multiple files and restores from multiple files, making it unnecessary to use the older **gsplit** command. Only SYSDBA and the owner of a database can back up a database. Any InterBase user defined on the server can restore a database, although the user must be SYSDBA or the database owner in order to restore it *over* an existing database.

Note When you back up and restore databases from IBConsole on Windows platforms, you are accessing this same tool through the IBConsole interface.

See Chapter 8, "Database Backup and Restore" for full documentation and reference on using **gbak**.

gfix

gfix configures several properties of a database, including:

- Database active/shutdown status
- Default cache allocation for clients
- Sweep interval and manual sweep
- Synchronous or asynchronous writes
- Detection of some types of database corruption
- Recovery of unresolved distributed transactions

You can also access all the functionality of **gfix** through the IBConsole graphical interface. Only SYSDBA and the owner of a database can run **gfix** against that database.

See Chapter 6, "Database Configuration and Maintenance" for descriptions of these properties, and a reference of the **gfix** tool.

gsec

You can configure authorized users to access InterBase servers and databases with **gsec**. You can also perform the same manipulations on the security database with IBConsole.

See Chapter 5, "Database Security" for full details and reference.

gstat

gstat displays some database statistics related to transaction inventory, data distribution within a database, and index efficiency. You can also view these statistics from IBConsole. You must be SYSDBA or the owner of a database to view its statistics.

See Chapter 9, "Database Statistics and Connection Monitoring" for more information on retrieving and interpreting database statistics.

iblockpr (gds_lock_print)

You can view statistics from the InterBase server lock manager to monitor lock request throughput and identify the cause of deadlocks in the rare case that there is a problem with the InterBase lock manager. The utility is called **gds_lock_print** on the UNIX platforms, and **iblockpr** on the Windows platforms.

See **Chapter 9, "Database Statistics and Connection Monitoring"** for more information on retrieving and interpreting lock statistics.

ibmgr

On UNIX servers, use the **ibmgr** utility to start and stop the InterBase server process. See the section **"Using ibmgr to start and stop the server"** on page 3-6 for details on using this utility.

Chapter

2

IBConsole: The InterBase Interface

InterBase provides an intuitive graphical user interface, called IBConsole, with which you can perform every task necessary to configure and maintain an InterBase server, to create and administer databases on the server, and to execute interactive SQL (ISQL). IBConsole enables you to:

- Manage server security
- Back up and restore a database
- View database and server statistics
- Perform database maintenance, including:
  - Validating the integrity of a database
  - Sweeping a database
  - Recovering transactions that are “in limbo”

IBConsole runs on Windows, but can manage databases on any server on the local network.

Starting IBConsole

To start IBConsole, choose IBConsole from the **Start** | **Programs** | **InterBase** menu. The IBConsole window opens:

Figure 2.1 IBConsole window

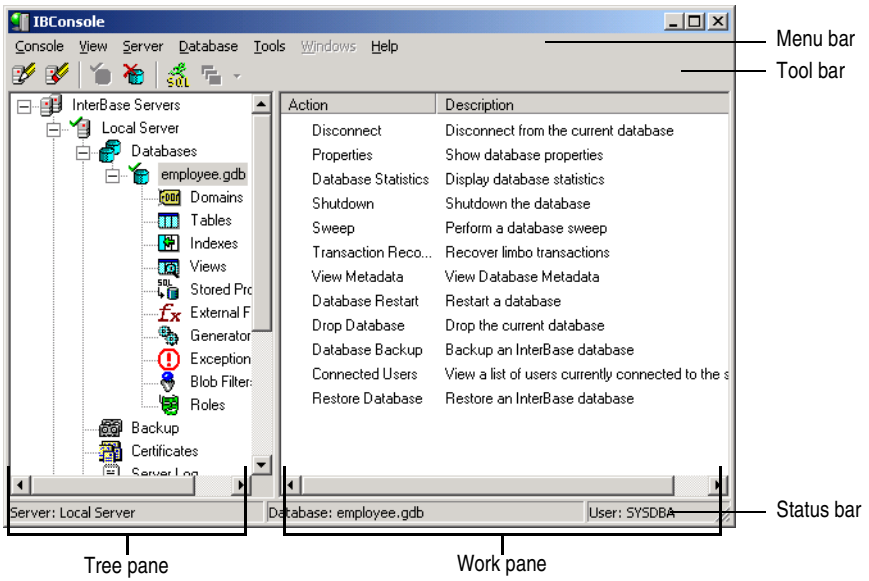

Elements in the IBConsole dialog:

- **Menu bar** Commands for performing database administration tasks.
- **Tool bar** Shortcut buttons for menu commands.
- **Tree pane** Displays a hierarchy of registered servers and databases.
- **Work pane** Displays specific information or allows you to perform activities, depending on what item is currently selected in the Tree pane.
- **Status bar** Shows the current server, user login, and selected database.

IBConsole menus

The IBConsole menus are the basic way to perform tasks with IBConsole. There are seven pull-down menus.

- **Console menu** enables you to exit from IBConsole.
- **View menu** enables you to indicate whether or not IBConsole displays system and temporary tables and dependencies and to change the display and appearance of items listed in the Work pane.

- **Server menu** enables you to register and un-register a server, log in to and log out of a server, diagnose a server connection, manage user security, add and remove certificates, view the server log file, and view server properties. For more information, see "Connecting to servers and databases" on page 4-2.
- **Database menu** enables you to register and un-register a database, create and drop a database, connect to and disconnect from a database, view database metadata, view a list of users connected to the database, view and set database properties, perform database maintenance, validation, and transaction recovery. For more information, see "Connecting to servers and databases" on page 4-2
- **Tools menu** enables you to add custom tools to the Tools menu and start the interactive SQL window. The interactive SQL window has its own set of menus, which are discussed in Chapter 10, "Interactive Query".
- **Windows menu** enables you to view a list of active IBConsole windows and to manage them. See "Switching between IBConsole windows" on page 2-7 for more information.
- **Help menu** enables you to access both IBConsole on-line help and InterBase on-line help.

Context menus

IBConsole also enables you to perform certain tasks with context sensitive popup menus called *context menus*. Tables 2.1 and 2.2 are examples of context menus.

When you right-click a server icon, a context menu is displayed listing actions that can be performed on the selected server.

Table 2.1 IBConsole context menu for a server icon

Popup command	Description
Register	Register the current server.
Un-register	Un-register the current server.
Login	Login to the selected server.
Logout	Logout from the current server.
Add Certificate	Add certificate ID/keys for the current server.
User Security	Authorize users on the current server.
View Logfile	Display the server log for the current server.
Diagnose Connection	Display database and network protocol communication diagnostics.
Properties	View and update server information for the current server.

When you right-click a connected database icon, a context menu is displayed listing actions that can be performed on the database:

Table 2.2 IBConsole context menu for a connected database icon

Popup command	Description
Disconnect	Disconnect from the current database.
Maintenance	Perform maintenance tasks including: view database statistics, shutdown, database restart, sweep, and transaction recovery.
Backup/Restore	Back up or restore a database to a device or file.
View Metadata	View the metadata for the selected database.
Connected Users	Displays a list of users connected to the database.
Properties	View database information, adjust the database sweep interval, set the SQL dialect and access mode, and enable forced writes.

IBConsole toolbar

A toolbar is a row of buttons that are shortcuts for menu commands. The following table describes each toolbar button in detail.

Figure 2.2 IBConsole Toolbar



Table 2.3 IBConsole standard toolbar

Button	Description
	**Register server**: opens the register server dialog, enabling you to register and login to a local or remote server. See "Registering a server" on page 4-2 for more information.
	**Un-register server**: enables you to unregister a local or remote server. This automatically disconnects a database on the server and logout from the server. See "Unregistering a server" on page 4-5 for more information.

Table 2.3 IBConsole standard toolbar (*continued*)

Button	Description
	**Database connect**: Connects to the highlighted database using the user name and password for the current server connection. See "Connecting to a database" on page 4-7 for more information.
	**Database disconnect**: enables you to disconnect a database on the current server. See "Disconnecting a database" on page 4-8 for more information.
	**Launch SQL**: opens the interactive SQL window, which is discussed in detail in Chapter 10, "Interactive Query".

Tree pane

When you open the IBConsole window, you must register and log in to a local or remote server and then register and connect to the server's databases to display the Tree pane. See "Connecting to servers and databases" on page 4-2 to learn how to register and connect servers and databases.

Figure 2.3 IBConsole Tree pane

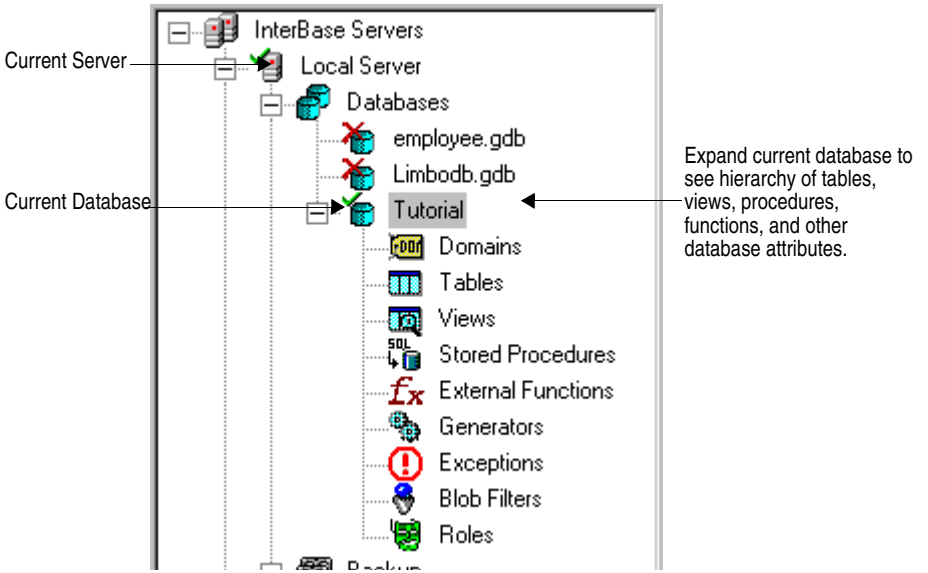

Navigating the server/database hierarchy is achieved by expanding and retracting nodes (or branches) that have subdetails or attributes. This is accomplished by a number of methods, outlined in Table 2.4.

To expand or retract the server/database tree in the Tree pane:

Table 2.4 Server/database tree commands

Tasks	Commands
Display a server's databases	• Left-click the plus (+) to the left of the server icon • Double-click the server icon • Press the plus (+) key • Press the right arrow key
Retract a server's databases	• Left-click the minus (–) to the left of the server icon • Double-click the server icon • Press the minus (–) key • Press the left arrow key

In an expanded tree, click a database name to highlight it. The highlighted database is the one on which IBConsole operates, referred to as the *current database*. The *current server* is the server on which the current database resides.

The hierarchy displayed in the Tree pane of Figure 2.3 is an example of a fully expanded tree.

- Expanding the InterBase Servers branch displays a list of registered servers.
- Expanding a connected server branch displays a list of server attributes, including Databases, Backup, Users, Certificates, and the Server Log.
- Clicking on the Database branch displays a list of registered databases on the current server.
- Clicking on Server Log displays the "View Logfile" action in the Work pane.
- Expanding a connected database branch displays a list of database attributes, including Domains, Tables, Views, Stored Procedures, External Functions, Generators, Exceptions, Blob Filters, and Roles.
- Expanding Backup shows a list of backup aliases.

Work pane

Depending on what item has been selected in the Tree pane, the Work pane gives specific information or enables you to execute certain tasks.

To display a list of backup aliases for the current server, click the Backup icon.

To display a list of InterBase certificate keys and IDs for the current server, click the Certificates icon.

To display a list of users defined on the server, click the Users icon.

To display information about a database attribute, click the database attribute icon.

To display information about a database object in a viewer, click the Work Pane icon for the object (for example, a table name).

See "Viewing metadata" on page 10-15 for more information.

Standard text display window

The standard text display window is used to monitor database backup and restoration, to display database statistics and to view server and administration logs.

The standard text display window contains a menu bar, a toolbar with icons for often-used menu commands, and a scrolling text display area. Figure 8.3, "Database backup verbose output," on page 8-7 is an example of the standard text display window.

Elements in a standard text display window:

- **Menu bar** The File menu enables you to save the contents of the window and Exit from the window. The Edit menu enables you to copy selected text in the window to the clipboard, select all text in the window, cut and paste text, and find a specified word or phrase within the displayed text.
- **Toolbar** Save and Copy toolbar buttons enable you to save the contents of the text display window as well as copy selected text to the clipboard.
- **Status bar** Shows the cursor location, given by line and column, within the text display window.

Switching between IBConsole windows

Use the Active Windows dialog to switch between IBConsole windows, or to close specific windows. To access the Active Windows dialog, click on the Windows menu. The dialog appears:

Figure 2.4 Active Windows dialog

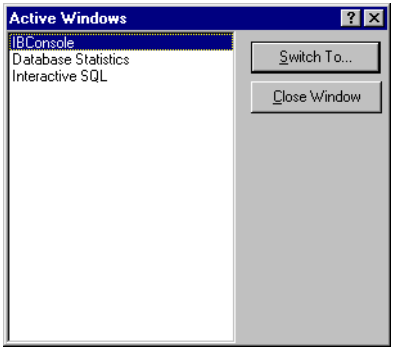

- To switch to a different IBConsole window, select it and click the Switch To button.
- To close a window, select it and click the Close window button.

Managing custom tools in IBConsole

Use the Tools dialog to add, edit, and delete custom tools for the IBConsole interface. To access the Tools dialog, select **Tools** | **Configure Tools** from the IBConsole menu. The Tools dialog is displayed:

Figure 2.5 Tools dialog

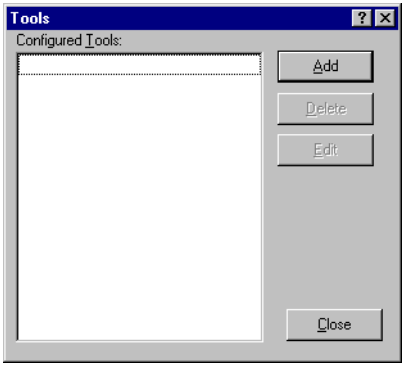

- To delete a tool, select it and click Delete.
- To modify a tool, select it and click Edit. Change the relevant fields in the Tool Properties dialog.
- To add a tool, click Add. The Tool Properties dialog appears.

Figure 2.6 Tool Properties dialog

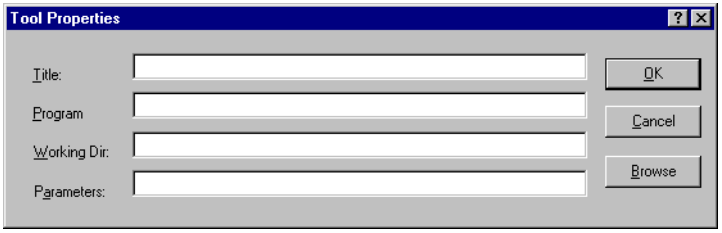

To add a custom tool:

- Enter the name of your utility in the Title field. This is the name that will be displayed on the Tools menu. Use an ampersand (&) to specify an accelerator key for the menu item. Conflicting accelerator keys are automatically resolved. If you do not specify an accelerator key, one will be chosen automatically.
- Enter the path and the executable to be launched in the Program field.
- Enter the working directory for your utility in the Working Dir field. If no working directory is specified, then it defaults to the current directory.
- Enter any other parameters needed to run your utility in the Parameters field.

Chapter

3

Server Configuration

This chapter describes the operation and configuration of the InterBase server process, including the following topics:

- Configuring server properties
- Using InterBase Manager to start and stop InterBase
- Starting and stopping the InterBase Server on UNIX
- Example initialization script installation on Linux
- Using environment variables
- Managing temporary files
- Configuring parameters in ibconfig
- Viewing the server log file

Configuring server properties

You can use InterBase Guardian to change database cache size of client map size. The InterBase Guardian Server Properties dialog enables you to display and configure these server settings. To access InterBase Guardian, right-click the InterBase Guardian icon in the System Tray. You can access the Server Properties dialog by any of the following methods:

- Select a server (or any branch under the server hierarchy) in the Tree pane and choose **Server | Server Properties**.
- Select a server in the Tree pane and click Server Properties in the Work pane.
- Right-click a server in the Tree pane and choose Server Properties from the context menu.

The Server Properties dialog contains two tabs, Alias and General.

The General tab

The General tab of the Server Properties dialog is where you can view such server settings as the version, capabilities, number of databases, and number of attachments. You cannot edit the information displayed on this tab.

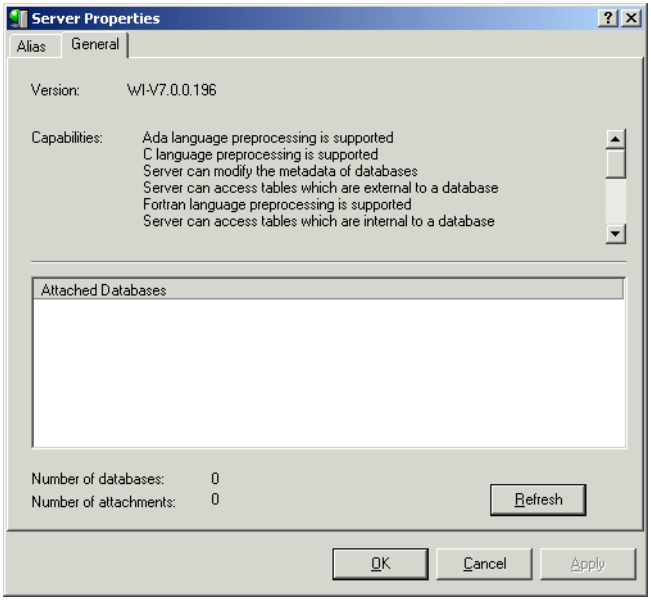

The server properties displayed are:

- **Version**: displays the version number for the InterBase Server.
- **Capabilities**: displays support capabilities for the InterBase Server.
- **Attached databases**: displays the path and filename for each attached database
- **Number of databases**: displays the total number of databases in the InterBase Server.
- **Number of attachments**: displays the total number attachments to the InterBase Server.

You cannot update the information displayed on the General tab; however, you can click Refresh at any time to retrieve the current server property information. If you need to view or configure server settings, click the IB Settings tab.

The Alias tab

On the Alias tab, you can inspect the host name and network protocol for the server. You can inspect and change the Alias name and description.

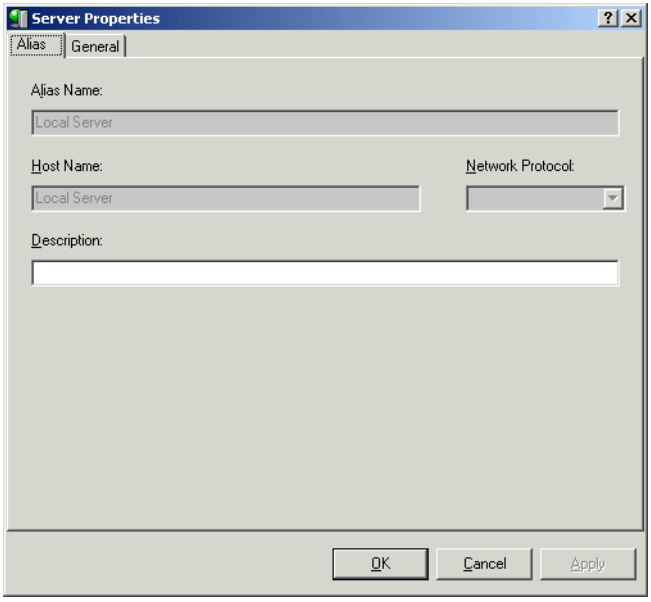

- **Alias Name**: the name of the server as it appears in the Tree pane. This setting is editable.
- **Host Name**: the name of the host server. This is determined at the time you create the server connection and cannot be changed in this dialog.
- **Network Protocol**: the protocol that the server is using to communicate. This is determined at the time you create the server connection and cannot be changed in this dialog.
- **Description**: any additional information you wish to add about the server. This field is optional and editable.

SMP support

InterBase provides symmetric multiprocessor (SMP) support for both clients and servers. Older versions of InterBase ran on SMP systems safely by allowing only a single processor at a time to execute within the InterBase components. Current versions of InterBase exploit SMP hardware by running InterBase threads on all processors simultaneously for increased throughput and performance.

Important When you purchase a single server license, you acquire the right to use a single processor. You must purchase one additional license for each additional processor that you wish to use.

On Windows platforms, the CPU_AFFINITY setting in the *ibconfig* configuration file specifies which processors of a multiprocessor system InterBase should use. The default setting, in effect when CPU_AFFINITY is commented out, is to use as many processors as licensing permits. See "**Expanded processor control: CPU_AFFINITY**" below for how to specify a subset of processors to use.

Expanded processor control: CPU_AFFINITY

On Windows multiprocessor platforms, you can specify which processors InterBase should use by adding the CPU_AFFINITY parameter to the *ibconfig* file. This setting is useful whenever the number of licensed processors is less than the number of actual processors present.

Important Note that when you purchase a single server license, you acquire the right to use a single processor. You must purchase one additional license for each additional processor that you wish to use.

The CPU_AFFINITY parameter populates a bit vector in which each bit represents a processor on the system on which the threads are allowed to run. The maximum number of processors depends on the operating system. To specify which processors should be used, give CPU_AFFINITY a decimal value that corresponds to the binary value that results from setting a bit for each desired machine. For example, to use processors 2 and 3, set CPU_AFFINITY to 6:

```
CPU_AFFINITY 6
```

To use these processors	CPU_AFFINITY value	Array setting
1	1	0 0 1
2	2	0 1 0
1 and 2	3	0 1 1
3	4	1 0 0
2 and 3	6	1 1 0
1, 2, and 3	7	1 1 1

ibconfig parameter: MAX_THREADS

Setting the MAX_THREADS parameter in *ibconfig* controls the maximum number of threads that can be active at one time within the InterBase engine. The default setting is 100:

```
MAX_THREADS 100
```

This configuration parameter allows the user to control the number of active threads in the engine. The ideal setting for this number depends partly on the nature of the work being performed by your clients. If you have many clients

performing very similar tasks, you may want to lower the MAX_THREADS setting to reduce contention. On the other hand, if simultaneous activity is highly diverse, setting this to a higher value may increase throughput.

Note that this setting does not affect the maximum possible threads that can be created by the InterBase server but only the number that can be active in the engine at one time.

Using InterBase Manager to start and stop InterBase

The InterBase Server and InterBase Guardian must be started before you enable database connections. On Windows platforms, you can use the InterBase Manager to start and stop the InterBase Server and Guardian. InterBase Manager is a Windows Control Panel applet. To start the InterBase Manager, choose **InterBase Manager** from the Control Panel. You can use InterBase Manager to do the following:

- Choose the server startup mode: whether to start the InterBase server manually, or have it start automatically at system boot time
- Change the path to the server: if you click the **Change** option, you can browse and select a different directory
- Change how InterBase Server operates. By default, InterBase runs automatically as a service on Windows NT or Windows 2000, though it is possible (but not recommended) to run it as an application. On Windows 98, InterBase runs only as an application.

Note To start InterBase Server as an application from a command prompt or in a batch file, invoke InterBase Guardian with option **-a**:

```
ibguard -a
```

InterBase Guardian starts the server, and places an icon in the System Tray.

- **Start** InterBase Server and InterBase Guardian, via a Start/Stop button. Click **Start** in the InterBase Manager Status area to start InterBase Server (and InterBase Guardian). The server status changes, and an InterBase Guardian icon appears in the system tray. Once you have started the InterBase Server, you can exit InterBase Manager, and both InterBase Server and InterBase Guardian will continue to run. The InterBase Guardian icon remains in the System Tray until you stop the server.
- **Stop** InterBase Server and InterBase Guardian, via a Start/Stop button. Click **Stop** in the InterBase Manager Status area to stop InterBase Server (and InterBase Guardian). Or, right-click the InterBase Guardian icon in the System Tray and choose **Shutdown**.

Starting and stopping the InterBase Server on UNIX

The following sections describe how to start and stop the InterBase server on UNIX.

Using ibmgr to start and stop the server

The InterBase Server process **ibserver** runs as a daemon on UNIX systems. Use **ibmgr** to start and stop the server process. To use **ibmgr**, you must be logged on to the server machine.

Syntax `ibmgr -command [-option [parameter] ...]`

or

```
ibmgr [enter key]
IBMGR> command [-option [parameter]]
```

Description On UNIX, the InterBase server process runs as a daemon. A daemon runs even when no user is logged in to the console or a terminal on the UNIX system.

ibmgr is a utility for managing the InterBase server process on UNIX systems. You must be logged on to the machine on which the server is running to use **ibmgr**.

Note The **ibmgr32.exe** file that is present in older Windows installations is an older client-side utility whose functions are entirely different than **ibmgr** on UNIX. The name is coincidental.

Options

start [-once | -forever]	Starts server; the **-forever** switch causes the server to restart if it crashes; default is **-forever**
shut	Rolls back current transactions, terminates client connections, and shuts down server immediately
show	Shows host and user
user *user_name*	Supplies SYSDBA
password *password*	Supplies SYSDBA password
help	Prints help text
quit	Quits prompt mode

Starting the server

To start the InterBase server, log in as the “root” or “interbase” user. (“interbas” is a synonym for “interbase,” to accommodate operating systems that do not support nine-character account names.) For example, start InterBase with the following command:

```
ibmgr -start
```

Note Once you have started **ibserver** using one login, such as "root," be aware that all objects created belong to that login. They are not accessible to you if you later start **ibserver** as one of the other two ("interbas" or "interbase"). It is highly recommended to run the InterBase Server as "interbase."

InterBase Classic uses the **inetd** process to handle incoming requests. There is no need to explicitly start the server; **inetd** forks off a process to handle incoming requests. Usually, **inetd** is set up to start automatically.

Stopping the server

Note For safety, make sure all databases have been disconnected before you stop the server.

The command switches **-user** and **-password** can be used as option switches for commands like **-start** or **-shut**. For example, you can shut down a server in any of the following ways:

```
ibmgr -shut -password password
```

or

```
ibmgr Enter
IBMGR> shut -password password
```

or

```
imbgr Enter
IBMGR> password password
IBMGR> shut
```

Note The **-shut** option rolls back all current transactions and shuts down the server immediately. If you need to allow clients a grace period to complete work and detach gracefully, use shutdown methods for individual databases. See "Shutting down and restarting databases" on page 6-30.

Starting the server automatically

To configure a UNIX server to start the InterBase Server automatically at server host boot-time, you must install a script that the *rc* initialization scripts can run. Refer to */etc/init.d/README* for more details on how UNIX runs scripts at boot-time.

Example initialization script

```
#!/bin/sh
# ibserver.sh script - Start/stop the InterBase daemon

# Set these environment variables if and only if they are not set.
: ${INTERBASE:=/usr/interbase}
# WARNING: in a real-world installation, you should not put the
```

```
# SYSDBA password in a publicly-readable file. To protect it:
# chmod 700 ibserver.sh; chown root ibserver.sh
export INTERBASE

ibserver_start() {
   # This example assumes the InterBase server is
   # being started as UNIX user 'interbase'.
   echo '$INTERBASE/bin/ibmgr -start -forever' | su interbase
}

ibserver_stop() {
   # No need to su.
   $INTERBASE/bin/ibmgr -shut -user SYSDBA -password password
}

case $1 in
'start') ibserver_start ;;
'start_msg') echo 'InterBase Server starting...\c' ;;

'stop') ibserver_stop ;;
'stop_msg') echo 'InterBase Server stopping...\c' ;;

*) echo 'Usage: $0 { start | stop }' ; exit 1 ;;
esac

exit 0
```

Example initialization script installation on Solaris

1 Log in as root.

```
$ su
```

2 Enter the example script above into the initialization script directory.

```
# vi /etc/init.d/ibserver.sh
```

3 Enter text

4 Link the initialization script into the *rc* directories for the appropriate run levels for starting and stopping the InterBase server.

```
# ln /etc/init.d/ibserver.sh /etc/rc0.d/K30ibserver
# ln /etc/init.d/ibserver.sh /etc/rc2.d/S85ibserver
```

Example initialization script installation on Linux

1 Log in as root.

```
$ su
```

2 Enter the Linux example script (given below) into the initialization script directory.

```
# cp ibserver.sh /etc/rc.d/init.d/ibserver.sh
# chmod 700 /etc/rc.d/init.d/ibserver.sh
```

3 Link the initialization script into the *rc* directories for the appropriate run levels for starting the InterBase server.

```
# ln -s /etc/rc.d/init.d/ibserver.sh /etc/rc.d/rc0.d/S85ibserver
```

4 Link the initialization script into the *rc* directories for the appropriate run levels for stopping the InterBase server.

```
# ln -s /etc/rc.d/init.d/ibserver.sh /etc/rc.d/rc0.d/K30ibserver
```

5 Make sure you have host equivalence:

```
# touch /etc/gds_hosts.equiv
# echo "+" >> /etc/gds_hosts.equiv
```

6 Make sure you don't have an **inetd** entry for InterBase Classic:

```
# echo -e "/gds_db/s/^/#/\nwq" | ed /etc/inetd.conf
# killall -HUP inetd
```

Example Linux initialization script

```
#!/bin/sh
# ibserver.sh script - Start/stop the InterBase daemon
# Set these environment variables if and only if they are not set.
: ${INTERBASE:=/usr/interbase}
# WARNING: in a real-world installation, you should not put the
# SYSDBA password in a publicly-readable file. To protect it:
# chmod 700 ibserver.sh; chown root ibserver.sh
export INTERBASE

ibserver_start() {
   # This example assumes the InterBase server is
   # being started as user "interbase".
   su - interbase -c "$INTERBASE/bin/ibmgr -start -forever"
   RETVAL=$?
   [ $RETVAL -eq 0 ] && touch /var/lock/subsys/ibserver
}

ibserver_stop() {
   # No need to su.
   $INTERBASE/bin/ibmgr -shut -user SYSDBA -password password
   RETVAL=$?
   [ $RETVAL -eq 0 ] && rm -f /var/lock/subsys/ibserver
}

if [ ! -d "$INTERBASE" ] ; then
   echo "$0: cannot find InterBase installed at $INTERBASE" >&2
   exit 1
fi
if [ ! -x "$INTERBASE/bin/ibmgr" ] ; then
   echo "$0: cannot find the InterBase SuperServer manager as
                $INTERBASE/bin/ibmgr" >&2
if [ ! -x "$INTERBASE/bin/gds_inet_server" ] ; then
```

```
   echo "$0: this is InterBase Classic; use inetd instead of
                ibserver daemon" >&2
   fi
   exit1
fi

case $1 in
'start')
   ibserver_start ;
   echo "InterBase started" ;;
'start_msg')
   echo 'InterBase Server starting...\c' ;;

'stop')
   ibserver_stop ;
   echo "InterBase stopped" ;;
'stop_msg')
   echo 'InterBase Server stopping...\c' ;;

'restart')
   ibserver_stop ; ibserver_start
   echo "InterBase restarted" ;;
'restart_msg')
   echo 'InterBase Server restarting...\c' ;;

*) echo "Usage: $0 { start | stop | restart }" ; exit 1 ;;
esac

exit 0
```

The attachment governor

The InterBase server has an attachment governor that regulates the number of attachments to the server. Multiply the value of the USERS field in the license file by four to determine the total number of permitted concurrent attachments.

All successful attempts to create or connect to a database increment the number of current attachments. Both local and remote connections count toward the connection limit. Connections are permitted by the governor until the maximum number of concurrent attachments is reached. All successful attempts to drop or disconnect from a database decrement the number of current attachments.

Once the maximum number of attachments is reached, the server returns the error constant *isc_max_att_exceeded* (defined in *ibase.h*), which corresponds to the message:

```
Maximum user count exceeded. Contact your database administrator.
```

Using environment variables

This section describes the environment variables that InterBase recognizes. When defining environment variables, keep these rules in mind:

- Environment variables must be in the scope of the **ibserver** process (SuperServer architecture) or **gds_inet_server** process (Classic architecture).
- On Windows NT/2000, define environment variables as *system variables* in the Windows **Control Panel | System** dialog.
- On UNIX, the easiest way to define environment variables is to add their definitions to the system-wide default shell profile.

ISC_USER and ISC_PASSWORD

If you do not provide a user name and password when you connect to a database or when you run utilities such as **gbak**, **gstat**, and **gfix**, InterBase looks to see if the ISC_USER and ISC_PASSWORD environment variables are set; if they are, InterBase uses that user and password as the InterBase user.

Although setting these environment variables is convenient, do not do so if security is at all an issue.

The INTERBASE environment variables

INTERBASE

The INTERBASE variable is used both during installation and during runtime. During installation, it defines the path where the InterBase product is installed. If this path is different from */usr/interbase*, all users must have the correct path set at runtime. During runtime, use the INTERBASE variable to set the InterBase install directory.

INTERBASE_TMP

The INTERBASE_TMP variable can be used to set the location of InterBase's sort files on the server. There are other options for defining the location of these files. See "Configuring sort files" on page 3-12.

INTERBASE_LOCK and INTERBASE_MSG

INTERBASE_LOCK sets the location of the InterBase lock file and INTERBASE_MSG sets the location of the InterBase message file. These two variables are independent of each other and can be set to different locations.

Important The environment variables must be in the scope of the **ibserver** process. On Windows platforms, define the variables as *system variables* in the **Control Panel | System | Environment** dialog. On UNIX, the easiest way to do this is to add the variable definition to the system-wide default shell profile.

The TMP environment variable

The TMP environment variable defines where InterBase stores temporary files, if the INTERBASE_TMP variable is not defined.

Managing temporary files

InterBase creates two types of temporary files: sort files and history list files.

The InterBase server creates sort files when the size of the internal sort buffer isn't big enough to perform the sort. Each request (for example, CONNECT or CREATE DATABASE) gets and shares the same list of temporary file directories. Each request creates its own temporary files (each has its own I/O file handle). Sort files are released when sort is finished or the request is released. If space runs out in a particular directory, InterBase creates a new temporary file in the next directory from the directory list. If there are no more entries in the directory list, it prints an error message and stops processing the current request.

The InterBase **isql** client creates the history list files to keep track of the input commands. Each instance creates its own temporary files, which can increase in size until they run out of disk space. Temporary file management is not synchronized between clients. When a client quits, it releases its temporary files.

Configuring history files

To set the location for history files, define the TMP environment variable on your client machine. This is the *only* way to define the location of history files. If you do not set the location for the history files by defining the TMP environment variable, an InterBase client uses whatever temporary directory it finds defined for the local system. If no temporary directory is defined, it uses */tmp* on a UNIX system or *C:\temp* on a Windows system.

Configuring sort files

You should make sure to have plenty of free space available for temporary sorting operations. The maximum amount of temporary space InterBase needs might be larger than the database itself in some cases.

Temporary sort files are always located on the server where the database is hosted; you should specify temporary directories on disk drives that are physically local to the server (not on mapped drives or network mounted file systems).

There are two ways to specify directories for sort files:

- You can add an entry to the *$INTERBASE/ibconfig* file to enable directory and space definition for sort files. The syntax is:

  ```
  TMP_DIRECTORY size "pathname"
  ```

Important The pathname **must** be in double quotes, or the config file will fail.

This defines the maximum size in bytes of each sort directory. You can list several directories, each on its own line with its own size specification and can specify a directory more than once with different size configurations. InterBase exhausts the space in each specification before proceeding to the next one.

For example, if you specify *dir1* with a size of 5,000,000 bytes, then specify *dir2* with 10,000,000 bytes, followed by *dir1* with 2,000,000 bytes, InterBase uses *dir1* until it reaches the 5,000,000 limit, then uses *dir2* until it has filled the 10,000,000 bytes allocated there, and then returns to *dir1* where it has another 2,000,000 bytes available. Below are the *ibconfig* entries that describe this configuration:

```
TMP_DIRECTORY 5000000 "C:\dir1"
TMP_DIRECTORY 10000000 "D:\dir2"
TMP_DIRECTORY 2000000 "C:\dir1"
```

- You can use the INTERBASE_TMP and TMP environment variables to define the location.

If you specify temporary directories in *ibconfig*, the server uses those values for the sort files and ignores the server environment variable values. If you don't specify configuration of temporary directories in *ibconfig*, then the server picks a location for a sort file based on the following algorithm:

1. Use the directory defined in INTERBASE_TMP environment variable
2. If INTERBASE_TMP isn't defined, use directory defined in TMP environment variable
3. If TMP isn't defined, default to the */tmp* directory (UNIX) or *C:\temp* (Windows)

Configuring parameters in *ibconfig*

You specify configuration parameters for InterBase Server by entering their names and values in the configuration file, *ibconfig*. Entries are in the form:

```
parameter <whitespace> value
```

- *"parameter"* is a string that contains no whitespace and names a property of the server being configured.
- *"value"* is a number or string that is the configuration of the specified property.

Note Each line in *ibconfig* is limited to 80 characters, including the parameter name and any whitespace.

The following is a summary of allowable entries in *ibconfig*:

Table 3.1 Contents of *ibconfig*

Parameter	Description
ADMIN_DB	• Name of the InterBase security database • Needed only if the security database is not named *admin.ib* • Must always be in the InterBase home directory
ANY_EVENT_MEM_SIZE	• Bytes of shared memory allocated for event manager • Default is 32768
ANY_LOCK_MEM_SIZE	• Bytes of shared memory allocated for lock manager • Default is 98304
ANY_LOCK_SEM_COUNT	• Number of semaphores for interprocess communication (Classic architecture only) • Default is 32
ANY_LOCK_SIGNAL	• UNIX signal to use for interprocess communication (Classic architecture only) • Default is 16
CPU_AFFINITY	• (Windows only) In an SMP system, sets which processors can be used • Default is all processors (entry is commented out)
CONNECTION_TIMEOUT	• Seconds to wait before concluding an attempt to connect has failed • Default is 180
DATABASE_CACHE_PAGES	• Server-wide default for the number of database pages to allocate in memory per database • Can be overridden by clients • See "Configuring the database cache" on page 6-24 for more information • Defaults: 2048 - SuperServer, 75 - Classic.
DEADLOCK_TIMEOUT	• Seconds before an ungranted lock causes a scan to check for deadlocks • Default is 10
DUMMY_PACKET_INTERVAL	• Seconds to wait on a silent client connection before the server sends dummy packets to request acknowledgment • Default is 60

Table 3.1 Contents of *ibconfig*

Parameter	Description
EXTERNAL_FILE_DIRECTORY	• If your external database files are not in <*interbase_home*>/ext, specify their location here. • For security reasons, do not put other files in this directory.
EXTERNAL_FUNCTION_DIRECTORY	• If your UDF library is not in <*interbase_home*>/UDF, specify its location here. • For security reasons, do not put other files in this directory.
LOCK_ACQUIRE_SPINS	• Number of spins during a busy wait on the lock table mutex • Relevant only on SMP machines • Default is 0
LOCK_HASH_SLOTS	• Tune lock hash list; more hash slots means shorter hash chains • Not necessary except under very high load • Prime number values are recommended • Default is 101
MAX_THREADS	• Controls the maximum number of threads that can be active at one time within the InterBase engine • Default is 100w
SERVER_CLIENT_MAPPING	• Size in bytes of one client's portion of the memory mapped file used for interprocess communication • Default is 4096
SERVER_PRIORITY_CLASS	• Priority of the InterBase service on Windows NT/2000/XP • The value 1 is low priority, 2 is high priority • Relevant on Windows NT/2000/XP only • Default is 1
SERVER_WORKING_SIZE_MAX	• Threshold above which the OS is requested to swap out all memory • Relevant only on Windows NT, 2000, and XP • Default is 0 (system-determined)
SERVER_WORKING_SIZE_MIN	• Threshold below which the OS is requested to swap out no memory • Relevant on Windows NT and Windows 2000 only • Default is 0 (system-determined)

Table 3.1 Contents of *ibconfig*

Parameter	Description
SWEEP_QUANTUM	• Maximum number of records that a garbage collector thread or a sweeper thread is allowed to work before yielding control back to the worker threads • If you increase SWEEP_QUANTUM, consider increasing SWEEP_YIELD_TIME also • Default is 10
SWEEP_YIELD_TIME	• Time in milliseconds the sweeper or garbage collector thread sleeps • Does not affect worker threads • If you increase SWEEP_YIELD_TIME, consider increasing SWEEP_QUANTUM also • Default is 1 millisecond
TMP_DIRECTORY	• Directory to use for storing temporary files • Specify number of bytes available in the directory and the path of the directory • List multiple entries one per line; each directory is used according to the order specified • Default is the value of the INTERBASE_TMP environment variable; otherwise */tmp* on UNIX or *C:\temp* on Windows NT/2000
USER_QUANTUM	• Maximum number of records that a worker thread (thread running an user query) is allowed to work before yielding control back to other threads • Default is 100
V4_EVENT_MEMSIZE	• Bytes of shared memory allocated for event manager • Default is 32768 • Overridden by ANY_EVENT_MEMSIZE
V4_LOCK_GRANT_ORDER	• 1 means locks are granted first come, first served • 0 means emulate InterBase V3.3 behavior, where locks are granted as soon as they are available; can result in lock request starvation • Default is 1
V4_LOCK_MEM_SIZE	• Bytes of shared memory allocated for lock manager • Default is 98304 • Overridden by ANY_LOCK_MEM_SIZE

Table 3.1 Contents of *ibconfig*

Parameter	Description
V4_LOCK_SEM_COUNT	• Number of semaphores for interprocess communication (Classic architecture only) • Default is 32 • Overridden by ANY_LOCK_SEM_COUNT
V4_LOCK_SIGNAL	• UNIX signal to use for interprocess communication (Classic architecture only) • Default is 16 • Overridden by ANY_LOCK_SIGNAL
V4_SOLARIS_STALL_VALUE	• Number of seconds a server process waits before retrying for the lock table mutex • Relevant on Solaris only • Default is 60

Viewing the server log file

InterBase Server logs diagnostic messages in the file *interbase.log* in the InterBase install directory. Any messages generated by **ibserver** are sent to this file. This can be an important source of diagnostic information if your server is having configuration problems.

Refer to the *Language Reference* for a list of error messages that can appear in this file.

IBConsole displays this log file in a standard text display window. To display the Server Log dialog:

- Select a server and expand it if it is not already expanded, click Server Log and then double-click **View Logfile** in the Work pane.
- Right-click a server in the Tree pane and choose **View Logfile** from the context menu.
- Select a server and then choose **View Logfile** from the **Server** menu.

Figure 3.1 Server Log dialog

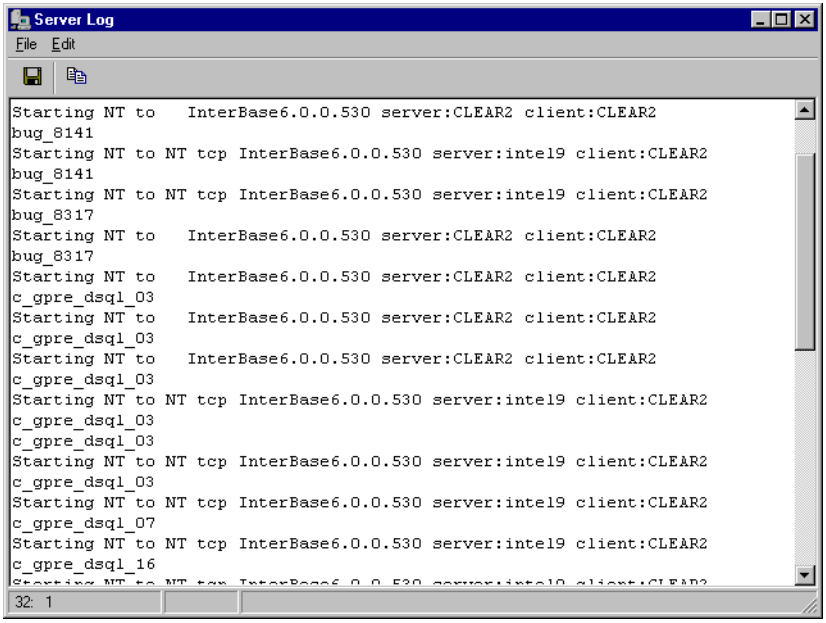

The standard text display window enables you to search for specific text, save the text to a file, and print the text. For an explanation of how to use the standard text display window, see "Standard text display window" on page 2-7.

Chapter

4

Network Configuration

This chapter details issues with configuring InterBase in a networked client/server environment. Topics include network protocols supported by InterBase, remote connection specifiers, and network troubleshooting tips.

Network protocols

InterBase supports TCP/IP for all combinations of client and server platforms. Additionally, InterBase supports NetBEUI for NT/2000/XP servers and Windows clients, and a local connection mode (involving interprocess communication but no network interface) for Windows clients.

InterBase is designed to allow clients running one operating system to access an InterBase server that is running on a different platform and operating system than the client. For example, a common arrangement is to have several Windows 98 PCs acting as client workstations concurrently accessing a departmental server running Windows 2000.

Table 4.1 Matrix of connection supported protocols

	Server platform		
Client platform	**Windows 98/ME**	**Windows NT/2000/XP**	**UNIX**
Windows 98/ME	TCP/IP, Local	TCP/IP, NetBEUI	TCP/IP
Windows NT/2000/XP	TCP/IP	TCP/IP, NetBEUI, Local	TCP/IP
UNIX/Linux	TCP/IP	TCP/IP	TCP/IP

Connecting to servers and databases

Before performing any database administration tasks, you must first register and log in to a server. Once you log in, you can register and connect to databases residing on the server. You can switch context from one connected database to another by selecting the desired database from the IBConsole Tree pane. The selected database in the Tree pane is referred to as the *current database*. The selected server or the server where the current database resides is referred to as the *current server*.

Registering a server

You can access the Register Server and Connect dialog in IBConsole by one of the following methods:

- Choose **Server | Register** or click the Register Server toolbar button.
- Double-click InterBase Servers in the Tree pane.
- Right-click InterBase Servers or any server in the Tree pane and choose Register from the context menu.

Figure 4.1 Register Server and Connect dialog

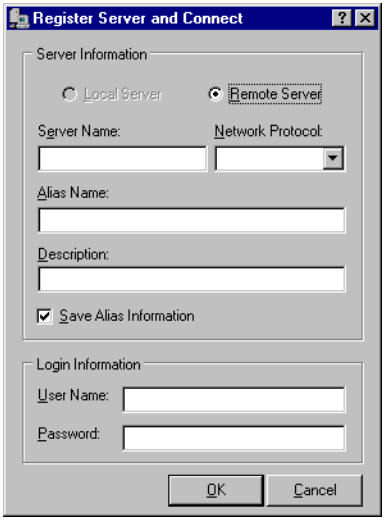

To register a local or remote server:

1. Select either the Local Server option or the Remote Server option.
2. If you choose Local Server, the Server Name, Network Protocol and Alias Name information is not required. These text fields are disabled. You can proceed to step 5.

3 If you choose Remote Server, type the name of the server in the Server Name text field, select a network protocol from the drop down list, and enter a server alias name in the Alias Name text field. Check the Save Alias Information check box if you wish to save the server alias name in the windows registry.

   Note The InterBase server name is the name of the database server machine. There is not a specific name for the InterBase server process itself. For example, if the server is running on the NT server "venus", you enter this name in the Server Name text field.

   The network protocol you select can be TCP/IP on any platform. On Windows, it can also be NetBEUI or local. Protocols are valid only when they are supported by both the client and the server.

4 Optionally, enter a description name for the server.

5 At this point you can choose to just register the server (without logging in) or you can choose to register and connect to the server simultaneously.

   - If you want to just register the server you can ignore the Login Information and click OK.
   - If you want to register and connect to the server simultaneously, enter a username and password in the corresponding text fields and click OK.

   Note The usernames and passwords must be the InterBase usernames and passwords stored in the InterBase security database (*admin.ib* by default) on the server.

Once a server is registered, IBConsole displays it in the Tree pane.

Logging in to a server

You can access the Server Login dialog in IBConsole by one of the following methods:

- In the Tree pane, select a registered server that is not already logged in. Choose **Server | Login** or double-click Login in the Work pane.
- In the Tree pane, double-click a registered server that is not already logged in.
- In the Tree pane, right-click a registered server that is not already logged in and choose Login from the context menu.

The Server Login dialog appears:

Figure 4.2 Server Login dialog

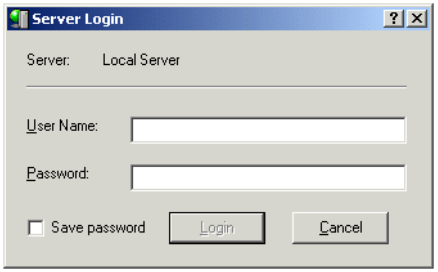

To log in to a server:

1. Verify that the server displayed in the Server field is correct.
2. Enter a username and password in the corresponding text fields. For convenience, IBConsole defaults the UserName text field to the last username that was used to log in (successfully or unsuccessfully).

   Note The usernames and passwords must be the InterBase usernames and passwords that are stored in the InterBase security database (*admin.ib* by default) on the server.

   The username is significant to 31 characters and is not case-sensitive. The password is significant to eight characters and is case-sensitive.

   All users must enter their username and password to log in to a server. The username and password are verified against records in the security database. If a matching record is found, the login succeeds.
3. Click Login to log in to the server.

Important Initially, every server has only one authorized user with username SYSDBA. The SYSDBA must log on and add other authorized users. For more information about how to add new users, see "User administration with IBConsole" on page 5-9.

Logging out from a server

You can log out from a server in IBConsole by one of the following methods:

- Select a connected server in the Tree pane (you can also select any branch under the desired server hierarchy) and choose **Server | Logout**.
- Select a connected server in the Tree pane and double-click Logout in the Work pane.
- Right-click a connected server in the Tree pane and choose Logout from the context menu.

A confirmation dialog asks you to confirm that you wish to close the connection to the selected server. Click Yes if you want to log out from the server, otherwise click No.

Note Logging out from a server automatically disconnects all databases but does not un-register any databases on the server.

Unregistering a server

You can unregister a disconnected server in IBConsole by one of the following methods:

- Select a server in the Tree pane and choose **Server | Un-register** or click the Unregister Server toolbar button
- Select a server in the Tree pane and double-click Un-register Server in the Work pane.
- Right-click a server in the Tree pane and choose Un-register from the context menu.

A confirmation dialog asks you to confirm that you wish to un-register the selected server. Click Yes if you want to un-register the server, otherwise click No.

Note Un-registering a server removes that server from the Tree pane and automatically logs you out of the current server as well as disconnects and un-registers any databases on the server.

Registering a database

You can access the Register Database and Connect dialog in IBConsole by one of the following methods:

- Choose **Database | Register**.
- Expand a connected server branch. Right-click Databases in the Tree pane and choose Register from the context menu.
- Select a disconnected database in the Tree pane and double-click Register in the work pane, or right-click the database and choose Register from the context menu.

The Register Database and Connect dialog appears:

Figure 4.3 Register Database and Connect dialog

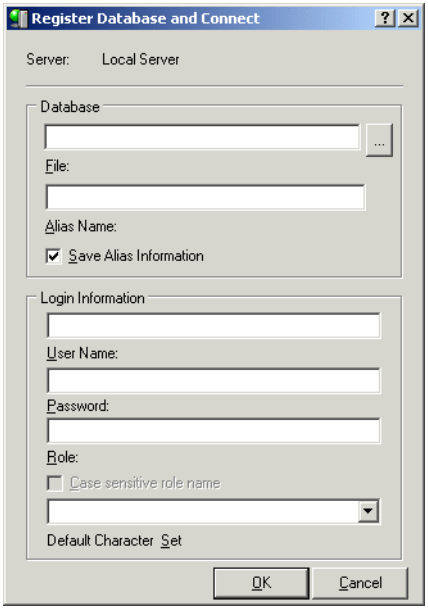

To register a database:

1. Make sure the server displayed in the Server field is correct.
2. Enter the database filename, including the path where the file is located, in the File text field. For databases that reside on the local server, you also have the option of clicking the Browse button ... to locate the file you want. The Browse button is disabled for all remote servers.
3. Type an alias name for the database in the Alias Name text field. This is the name that will appear in the IBConsole window. If you omit this step, the alias defaults to the filename that you select in step 2.
4. Check the Save Alias Information check box if you wish to permanently register the database. This saves the database alias name in the Windows registry.
5. At this point you can choose to just register the database without connecting, or you can choose to register and connect to the database simultaneously.

   If you only want to register the database, ignore the Login Information and click OK.
6. If you want to register and connect a database simultaneously, type the username, password and optional role and default character set for the database in the corresponding text fields and click OK.

Note If you want to connect using a role, specify the role in the Role text field. This is optional. Connecting using a role gives you all privileges that have been assigned to that role, assuming that you have previously been granted that role with the GRANT statement. For more information on roles, refer to "SQL roles" on page 5-6.

Once you register a database it appears in the Tree pane.

Connecting to a database

IBConsole provides two methods for connecting to a database. The first method is a quick connect using the username and password that were supplied with the login to the server to instantaneously connect the database. The second method allows you to connect to the database using a different username and password by accessing the Database Connect dialog.

Connect

If you want to perform an automatic connect, using the username and password supplied for the server login to instantaneously connect the database, you can do so by one of the following methods:

- Select a disconnected database in the Tree pane. Choose **Database | Connect**, choose Connect in the Work pane, or click on the Database Connect toolbar button.
- Right-click a disconnected database in the Tree pane and choose Connect from the context menu.
- Double-click a disconnected database in the Tree pane.

Once you connect to a database, the database tree expands to display the database hierarchy.

Connect as

If you want to access the Connect Database dialog in IBConsole to connect to the database using a different username and password from that which was supplied in the server login you can do so by one of the following methods:

- Select a disconnected database in the Tree pane. Choose **Database | Connect As** or choose Connect As in the Work pane.
- Right-click a disconnected database in the Tree pane and choose Connect As from the context menu.

The Database Connect dialog appears:

Figure 4.4 Database connect dialog

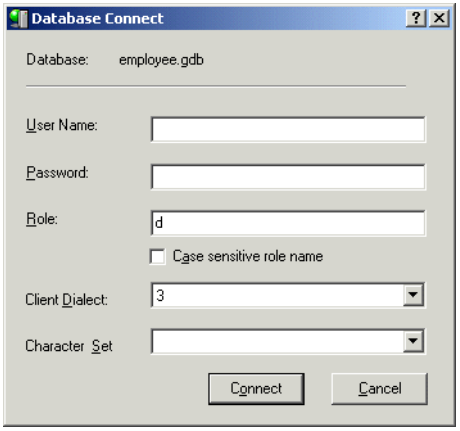

To connect to a database:

1. Verify that the database displayed in the Database field is correct.
2. Type the username and password for the database in the corresponding User Name and Password text fields.
3. If you want to connect as a role, specify the role in the Role text field. This is optional. Connecting as a role gives you all privileges that have been assigned to that role, assuming that you have previously been granted that role with the GRANT statement. Once you have typed a character in the Role field, the Case Sensitive Role Name field becomes active. Check this box if you want the server to consider case in the role name. Role names are case insensitive by default.

   For more information on roles, refer to "SQL roles" on page 5-6
4. Select the SQL Client dialect. The dialect for the database connection will default to the lower value of the client or server. For more information on SQL dialects, refer to "Understanding SQL Dialects" in the migration appendix of the InterBase *Operations Guide*.
5. Optionally, you can choose a character set to use. If you do not specify one here, the server uses the default set that was specified at creation time.
6. Click Connect.

Once you connect to a database, the database tree expands to display the database hierarchy.

Disconnecting a database

You can disconnect a database in IBConsole by one of the following methods:

- Select a connected database in the Tree pane (you can also select any branch under the desired database hierarchy) and choose **Database | Disconnect** or click the Disconnect Database toolbar button

- Select a connected database in the Tree pane and double-click Disconnect in the Work pane.
- Right-click a connected database in the Tree pane and choose Disconnect from the context menu.

A confirmation dialog asks you to confirm that you wish to close the connection to the selected database. Click OK if you want to disconnect the database, otherwise click Cancel.

Unregistering a database

You can unregister a disconnected database in IBConsole by one of the following methods:

- Select a database in the Tree pane (you can also select any branch under the desired database hierarchy) and choose **Database | Un-register**.
- Select a database in the Tree pane and double-click Un-register in the Work pane.
- Right-click a database in the Tree pane and choose Un-register from the context menu.

A confirmation dialog asks you to confirm that you wish to un-register the database. Click Yes if you want to un-register the database, otherwise click No.

Note Un-registering a database automatically disconnects the current database and removes it from the Tree pane.

Connection-specific examples

Here are some examples of connecting to databases on various types of servers.

- For a Windows server, the database path name must contain the appropriate drive letter designation. For example, to connect to a local database:

```
D:\users\accting\fin\accred.ib
```

- To connect to a database on a remote server using the TCP/IP protocol:

```
ntserver:D:\users\accting\fin\accred.ib
```

- To connect via NetBEUI (Windows NT/2000 servers only), use UNC notation:

```
\\ntserver\D:\users\accting\fin\accred.ib
```

- For a UNIX or Linux server, you must enter the complete and absolute directory path for the database. For example:

```
server:/usr/accting/fin/accred.ib
```

Connection troubleshooting

This section describes some troubleshooting guidelines for issues related to network configuration and client/server connections. If you are having trouble connecting client to server over a network, use the steps listed below to diagnose the cause. On Windows, you can perform some of these tests using the Communications Diagnostic dialog. See "Communication diagnostics" on page 4-15 for more information.

Connection refused errors

If the client fails to reach the server host at all, or the *gds_db* service fails to answer, you might get a "connection refused" error. Below is a checklist that you can use to diagnose the source of this error.

Is there low-level network access between the client and server?

You can quickly test whether the client cannot reach the server because of a physically disconnected network or improper network software configuration, by using the **ping** command. Usage is:

```
ping servername
```

Error messages from **ping** indicate that there is a network problem. Check that the network is plugged in, that the network wires are not damaged, and that the client and server software is properly configured.

Test connectivity from the client in question to another server; if it succeeds, this could rule out improper network configuration on the client.

Test connectivity from another client to the InterBase server host; if it succeeds, this could rule out improper network configuration on the server.

Can the client resolve the server's hostname?

InterBase clients must specify the server by name, not by IP address, except in some Linux distributions. Therefore, the client must be able to resolve the server's hostname. For TCP/IP, this is done either by maintaining a *hosts* file on the client with the mappings of hostnames to IP addresses, or by the client querying a DNS server or WINS server to resolve this mapping. Make sure the name server has a correct entry for the server host in question.

Is the server behind a firewall?

If the database server is behind a software or hardware firewall, all network traffic could be restricted and the client might not be able to reach the server at all. Some firewalls permit or restrict traffic based on the port to which the client attempts to connect. Because of this, it is not conclusive whether a given service can reach the server. Neither is it an indication of connectivity if the client can resolve the IP address; that merely indicates that the client can reach a name server that resolves the InterBase server host's name.

If the client is separated from the server by a firewall, the client cannot connect.

Are the client and server on different subnets?

NetBEUI cannot route network traffic between subnets. Other protocols can also be configured to restrict traffic between subnets. If the client and server are on a complex network with multiple subnets, ask your network administrator if the network configuration allows you to route network traffic between the client and server in question using a given protocol.

Can you connect to a database locally?

To confirm that the **ibserver** process is running on the server and able to attach to your database, try a local database connection:

1. Log in to the console of the database server host, and run an application such as command-line **isql**.

2. Attempt to connect to a database without specifying a hostname: list just the path.

The Communications Diagnostic dialog also has a local database attachment test. See "DB Connection tab" on page 4-16 for details.

Note Local connection mode is not available on UNIX servers.

Can you connect to a database loopback?

You can simulate a client/server connection and test the server's configuration without the additional variable of the client configuration and intervening network by connecting in a *loopback* mode.

1. Log in to the console of the database server host and run an application such as command-line **isql** or InterBase IBConsole ISQL.

2. Attempt to connect to the database using a remote connection specification, even though the server named is also the client host.

Whether this test fails or succeeds, it helps to narrow the focus of further diagnostic tests. If it fails, you can infer that the server's configuration is at fault. If it succeeds, you can infer that the server is not at fault and you can concentrate further tests on the client.

Is the server listening on the InterBase port?

If the **ibserver** process on the server has not started, there is no answer to attempts to connect to the *gds_db* service (port 3050).

Start the **ibserver** process on the server. Use **ibmgr -start** on UNIX, or the InterBase Manager on Windows. See Chapter 3, "Server Configuration."

Is the services file configured on client and server?

The *services* file must have correct entries to indicate the port number associated with the named service *gds_db*. This configuration must be accessible on the client as well as the server.

```
gds_db      3050/tcp      # InterBase Server
```

On Windows NT/2000, this file is in *C:\windows\system32\drivers\etc\services*.
On Windows 98/ME, this file is in *C:\windows\services*.
On UNIX, this file is in */etc/services*.

In a UNIX environment with NIS, the NIS server can be configured to supply the *services* file to all NIS clients on UNIX workstations.

Is the UNIX inetd daemon configured for InterBase Classic architecture?

When running the InterBase Superserver architecture, you should check the */etc/inetd.conf* file to make sure that the **inetd** daemon is not configured to listen on the **gds_db** service (port 3050). In InterBase SuperServer, the **ibserver** process takes over the task of listening on the port, and if both **inetd** and **ibserver** attempt to listen, then there is a conflict and the result is that neither can successfully accept connection requests.

Make the following change:

- Use a text editor to remove the line in */etc/inetd.conf* that mentions the **gds_db** service
- Restart **inetd** by sending it a HUP signal

The installation script in InterBase for UNIX is supposed to perform this task, but if something goes wrong, or the */etc/inetd.conf* file is restored to its configuration for InterBase Classic, you need to correct the configuration.

Connection rejected errors

If the client reaches the server host and the *gds_db* service answers but you still cannot attach to a database, it can result in a "connection rejected" error. Below is a checklist that you can use to diagnose the source of this error.

Did you get the correct path to the database?

Verify that you supplied the correct path to the database file. Keep in mind:

- On NT/2000, you must supply the drive letter with the path.
- On UNIX, paths are case-sensitive.
- Slash ("/") vs. backslash ("\") does not matter, unless you need to use double-backslashes in string literals in C or C++ code.

Is UNIX host equivalence established?

To use the UNIX user-equivalence feature, there must be a *trusted host* relationship between the client and the server. See "Users on UNIX" on page 5-2.

Is the database on a networked file system?

A database file must not reside on an NFS file system or a mapped drive. When the **ibserver** process finds such a case, it either denies the connection or passes the connection request on to the InterBase service running on the file server. See "Networked file systems" on page 6-6 for more details.

To correct this situation, move your database to a file system on a hard disk that is physically local to the database server.

Are the user and password valid?

The client application must use a valid user and password combination that matches an entry in the InterBase security database (*admin.ib* by default).

Does the server have permissions on the database file?

The **ibserver** process must have permission to read and write the database file at the operating system level. Check the permissions on the database file, and the uid of the **ibserver** process. (On UNIX, you have the option of running **ibserver** as user *interbase*, a non-superuser uid.)

The the InterBase security database (*admin.ib* by default) that contains users and passwords must also be writable by the **ibserver** process.

Does the server have permissions to create files in the InterBase install directory?

The **ibserver** process must have write permission in the InterBase directory (by default, */usr/interbase* on UNIX, *C:\Program Files\Borland\InterBase* on Windows). The server process must be able to write to, and perhaps create, the *interbase.log* file and other temporary files.

Disabling automatic Internet dialup

Microsoft Windows operating systems offer a networking feature that is convenient for users who use a modem to connect to the Internet: any TCP/IP request that occurs on the system activates an automatic modem dialing program. This is helpful for users who want to connect quickly as they launch a web browser or email client application.

This convenience feature is unnecessary on systems that use a client/server application to access an InterBase server on a local network. The TCP/IP service request that the client invokes triggers the Windows automatic modem dialer. This interferes with quick network connections from client to server.

This section describes several methods to suppress the automatic modem dial feature of Windows operating systems. No more than one of these methods should be necessary to accomplish the networking configuration you need.

Reorder network adapter bindings

You probably have a dialup adapter and an ethernet adapter for your local network. On Windows NT/2000, you can reverse the bindings order for your two adapters to force the ethernet adapter service the TCP/IP request before the dialup adapter tries. You can do this in **Control Panel | Networking | Bindings | All Adapters | Move Down**.

The local ethernet adapter satisfies TCP/IP requests it can, and those requests that can't be done locally—such as Internet requests—are passed on to the next adapter in the list, the dialup adapter.

Disabling autodial in the registry

Perform the following:

1. Start the registry editor with **regedit.exe**
2. Move to the registry key `HKEY_LOCAL_MACHINE\Software\Microsoft\ Windows\ CurrentVersion\Internet Settings: EnableAutoDial`
3. Change the value from 0 to 1

Preventing RAS from dialing out for local network activity

Perform the following if you are using Windows NT RAS:

1. Start the registry editor, with **regedit.exe**
2. Move to the registry key `HKEY_CURRENT_USER\Software\Microsoft\RAS Autodial\ Addresses`

   A better way to view these is to type **rasautou -s** from the command prompt
3. In the subkeys look for the local address and name; select the key and select **Delete** from the **Edit** menu
4. Close the registry editor

You might also wish to add addresses to the disabled list:

5. Start the registry editor with **regedt32.exe**, not **regedit.exe**
6. Move to the registry key `HKEY_CURRENT_USER\Software\Microsoft\RAS Autodial\Control`
7. Double click Disabled Addresses and add the address on a new line; click **OK** when you are finished
8. Close the registry editor

You must reboot the machine in both of the above cases.

Other errors

Unknown Win32 error 10061

This error is often associated with a missing server-access license for the InterBase software on the server host. Make sure you have licensed InterBase server to allow clients to connect from the network.

Unable to complete network request to host

This error occurs in cases when the InterBase client cannot establish a network connection to the server host. This can occur for a variety of reasons. Below is a list of common causes:

- The BDE Administrator requires that you specify the InterBase connect string in the `SERVER NAME` alias property. You must use this property and must not use the `PATH` alias property, or else you receive the network error message.

- The InterBase client attempts to translate the server portion of your connect string to an IP address, by calling `gethostbyname()`. If you supplied an IP address, `gethostbyname()` is likely to fail to resolve it. Some modern TCP/IP drivers— including Winsock 2 and Linux TCP/IP—do resolve strings that look like IP addresses. If you are on Windows, specify hosts by name, or else upgrade your TCP/IP driver to Winsock 2.
- The InterBase client must look up the InterBase network *service* by name. If the client doesn't find the entry for **gds_db** in the *services* file, it might fail to connect to the server, and give the network error. You can create the entry in the *services* file manually, or reinstall InterBase to perform this task.
- The server you specify must be running on the network that you use. If the hostname corresponds to a host that is inaccessible because of network interruption, or the host is not running, then the connection request fails with the network error.
- The syntax of the InterBase connect string determines the network protocol the client uses to connect to the server host (see "Connection-specific examples" on page 4-9). Different server platforms support different subsets of network protocols. If your server does not support the protocol indicated by your connect string, the connection attempt fails with the network error. For example, the NetBEUI connection syntax (*server*\C:*path**database.ib*) works only if your server is a windows NT, 2000, or XP server. The syntax does not work if your server is running UNIX or Linux.
- A network connection request succeeds only if the InterBase server is installed and active on the server host, and that the InterBase server is licensed to receive remote connection requests. If there is no process listening for connection requests, the client's connection requests with the network error. You should check that the InterBase server is installed on the server, that it is running, and that the license includes the Server capability.

Communication diagnostics

Network configuration of a client/server system involves several different software and hardware layers and proper configuration of each of these layers. When one or more layers are misconfigured, it is not always evident where the problem lies. InterBase Communication diagnostics helps to identify the source of the problem by testing each layer progressively for existing or potential network problems.

You can access the Communication Diagnostics dialog by one of the following methods:

- Select a disconnected server in the Tree pane. Choose **Server | Diagnose Connection.**

- Right-click InterBase Servers or any disconnected server in the Tree pane and choose Diagnose Connection from the context menu.
- Select a disconnected server from the Tree pane and double-click Diagnose Connection in the Work pane.

There are four types of diagnostics that you can perform. The Communications Diagnostics dialog has separate tabs for each diagnostic type.

DB Connection tab

This test lets you connect to an InterBase database using the InterBase client libraries. It is the most basic test of InterBase operation and is generally used only after confirmation that the underlying network is working correctly.

Figure 4.5 Communications dialog - DB Connection

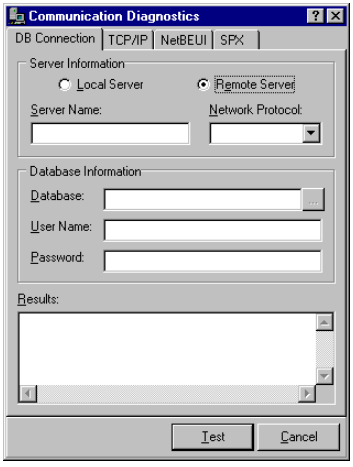

To run a DB Connection test:

1 Select either the Local Server option or the Remote Server option.

2 If you choose Local Server, the Server Name and Network Protocol information is not required. These text fields are disabled. You can proceed to step 5.

3 If you choose Remote Server, type the name of the server in the Server Name text field.

Note The InterBase server name is the name of the database server machine. There is not a specific name for the InterBase server process itself. For example, if the server is running on the NT server "*venus*", you enter this name in the Server Name text field.

4 If you choose Remote Server, select a network protocol from the drop down list: either TCP/IP, NetBEUI, named pipe, or local. Protocols are valid only when they are supported by both the client and the server.

5 Enter the database filename, including the path where file is located, in the Database text field. If you selected the Local Server option in step 1 you can also click the browse button ... to locate the file you want. If you selected the Remote Server option, however the browse button is disabled.

6 Type the username and password for the database in the corresponding User Name and Password text fields.

7 Click Test to display the results of the connectivity test in the Results text area.

Sample output (local connection)

```
Attempting to attach to:
   C:\Program Files\Borland\InterBase\examples\Database\employee.gdb
   Attaching ...Passed!
   Detaching ...Passed!

InterBase Communication Test Passed!
```

TCP/IP tab

Use this property sheet to test Winsock TCP/IP connectivity.

Figure 4.6 Communications dialog - TCP/IP

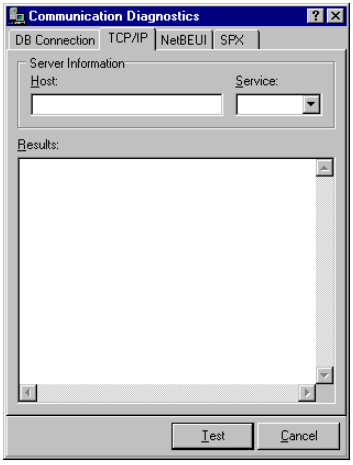

To run a winsock TCP/IP connectivity test:

1 Enter either a network host name or IP address in the Host text field.

2 Select a service name or number from the dropdown Service list. Possible service selections are: 21, Ping, 3050, ftp, **gds_db**.

Select Ping from the Service dropdown list to display a summary of round-trip times and packet loss statistics.

3 Click Test to display the results of the connectivity test in the Results text area.

Sample results (ftp):

```
Initialized Winsock.

Attempting connection to DBSERVE.
Socket for connection obtained.

Found service 'FTP' at port '21'.
Connection established to host 'DBSERVE' on port 21.

TCP/IP Communication Test Passed!
```

Sample results (ping):

```
Pinging DBSERVE [200.34.4.5] with 32 bytes of data.

Reply from 200.34.4.5: bytes=32 time=1ms TTL=128
Reply from 200.34.4.5: bytes=32 time=1ms TTL=128
Reply from 200.34.4.5: bytes=32 time=1ms TTL=128
Reply from 200.34.4.5: bytes=32 time=0ms TTL=128

Ping statistics for 200.34.4.5:
    Packets: Send = 4, Received = 4, Lost = 0 (0%),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

Table 4.2 Using Communication Diagnostics to diagnose connection problems

If the error message is...	Then check...
Failed to find named port	• Your *services* file to be sure there is an entry for *gds_db* in the form: **gds_db 3050/tcp**
Failed to connect to host	• Hostname, port 3050 • The InterBase Server to make sure it is installed properly, is running, and is configured for TCP/IP
Failed to resolve hostname	• Hostname • Your *hosts* file or DNS to be sure it has an entry for the server • That you used a hostname and not an IP address
Unavailable database	• Whether the InterBase server is running; the server must be running before attempting a database connection

NetBEUI tab

Use this property sheet to test NetBEUI connectivity between the client and the server.

Figure 4.7 Communications dialog - NetBEUI

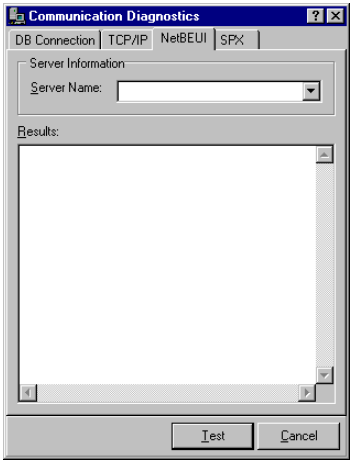

To run a NetBEUI connectivity test:

1. Select a Windows NT or Windows 2000 server on which InterBase has been installed from the Server Name drop down list. If the desired server does not exist in this list, you can type the server name in the edit portion of the drop down list.
2. Click Test to display the results of the connectivity test in the Results text area.

Sample output (NetBEUI connection):

```
Attempting to attach to DBSERVE using
the following named pipe:
  \\dbserve\pipe\interbas\server\qds.db.

NetBEUI Communication Test Passed!
```

Note NetBEUI is supported on all Windows clients, but only Windows NT and Windows 2000 support NetBEUI as a server.

The connection may fail if a Microsoft Windows network is not the default network for the client. You should also be logged into the MS Windows network with a valid NT/2000 user name and password.

Chapter

5

Database Security

InterBase provides several methods for configuring and enforcing security by controlling how a database is accessed and used. Server security enables you to:

- Add a user to the security database
- Delete a user from the security database
- Modify user information in the security database
- Display a list of users in the security database

This chapter gives an overview of these options. The user administration tools are covered here, but SQL statements for configuring privileges are in other InterBase books; these passages are referenced where appropriate.

Security model

Security for InterBase relies on a central security database for each server host. This database, *admin.ib* by default, contains a record for each legitimate user who has permission to connect to databases and InterBase services on that host. Each record includes the user login name and the associated encrypted password. The entries in this security database apply to all databases on that server host.

The username is significant to 31 characters and is not case sensitive. Password is significant to eight characters and is case sensitive.

Before performing any database administration tasks, you must first log in to a server. Once you log in to a server, you can then connect to databases residing on the server.

All users must enter their username and password to log in to a server. The password is encrypted for transmission over the network. The username and password are verified against records in the security database. If a matching record is found, the login succeeds.

The SYSDBA user

Every InterBase server has a SYSDBA user, with default password *masterkey*. SYSDBA is a special user account that can bypass normal SQL security and perform tasks such as database backups and shutdowns.

Initially, SYSDBA is the only authorized user on a server; the SYSDBA must authorize all other users on the server. Only the SYSDBA user can update the security database to add, delete, or modify user configurations. SYSDBA can use either **gsec** or IBConsole to authorize a new user by assigning a username and password in the security database.

Important We *strongly* recommend you change the password for SYSDBA as soon as possible after installing InterBase. If you do not alter the SYSDBA password, unauthorized users have easy access and none of your databases are secure.

Other users

The SYSDBA account can create other users on a per-server basis. Use **gsec** or IBConsole to create, modify, or remove users from the InterBase security database. These users are authorized to connect to any database on that database server host. It is a common design strategy to create a distinct InterBase user for each person who uses the databases on your server. However, other strategies are also legitimate. For example:

- Create one InterBase user for an entire group of people to use, in order to simplify password administration. For example, a user FINANCE could satisfy the access needs for any and all staff in a financial analysis team. This team only needs to remember one password between them.
- Create one InterBase user for a group of people to use, as warranted by requirements of distinct privilege configurations. For example, if Erin and Manuel have identical access to the data within a database, they could use the same InterBase user account.

Users on UNIX

If both the client and the server are running UNIX, you can allow UNIX usernames access to databases by configuring the server host to treat the client host as a *trusted host*.

To establish a trusted host relationship between two hosts, add an entry in */etc/hosts.equiv* or */etc/gds_hosts.equiv* on the server. The former file establishes trusted host status for any service (for example, **rlogin**, **rsh**, and **rcp**); the latter file

establishes trusted host status for InterBase client/server connections only. The format of entries in both files is identical; see your operating system documentation on *hosts.equiv* for details.

The login of the client user must exist on the server. In addition to the *hosts.equiv* method of establishing a trusted host, the you can also use the *.rhosts* file in the home directory of the account on the server that matches the account on the client.

The InterBase client library defaults to using the current client's UNIX login as the InterBase login only when the client specifies no username through any of the following methods:

- Database parameter buffer (*dpb*) parameters—see the *API Guide*
- Command-line options—for example, **-user** options of **isql** or another utility
- Environment variables—see "ISC_USER and ISC_PASSWORD" on page 3-11

Notes

- This feature is not implemented in InterBase 6 on a Windows NT or Windows 2000 server, because Windows does not implement a trusted host mechanism as UNIX does.
- Windows clients cannot be treated as trusted hosts by UNIX servers.

The InterBase security database

The InterBase server stores the names and passwords of its authorized users in a special security database that resides in the InterBase home directory. By default, it is named *admin.ib*.

You can use another name for the security database if you wish. If you change this name, you must add an entry to the *ibconfig* file, setting `ADMIN_DB` to the new name.

```
ADMIN_DB newname.ib
```

Note In older versions of InterBase, the security database was named *isc4.gdb*. Because files with a *gdb* extension automatically get backed up whenever they are touched in some versions of Windows XP and ME, using this extension degrades database performance. Therefore, InterBase recommends using a different extension for database names.

Every user of an InterBase server requires an entry in the InterBase security database. The **gsec** security utility lets you display, add, modify, or delete information in the security database. IBConsole provides a graphical interface for the same functionality. The following table describes the contents of the security database:

Table 5.1 Format of the InterBase security database

Column	Required?	Description
User name	Yes	The name that the user supplies when logging in; maximum length is 31 characters
Password	Yes	The user's password • Case sensitive • Only the first eight characters are significant • Maximum length: 32 characters.
UID	No	An integer that specifies a user ID
GID	No	An integer that specifies a group ID
Full name	No	User's real name (as opposed to login name)

System table security

InterBase stores the database metadata in its system tables. These tables have an intricate set of dependencies between them, and writing to one without sufficient knowledge can corrupt the database. For this reason, the system tables have the following *default* security access applied to them:

- By default, PUBLIC users have only SELECT privileges on the system tables.
- The database owner, the SYSDBA user, and the operating system administrator (root on UNIX and Administrator on Windows) have full access to the system tables, including write permission. These users can, if desired, assign write privileges to individual users or to PUBLIC, using the GRANT statement.

Older databases

InterBase applies this default security (no write access for PUBLIC) to older databases whenever possible:

- The **gbak** backup/restore utility applies the default security to any database when it is restored to ODS 10.1 (InterBase 6.5) or later.
- When an InterBase server that is version 6.5 or later attaches an older database, it applies the default privileges to that database if they are not already present, even if the database is ODS 10.0 or earlier.

Scripts for changing database security

Three SQL scripts are included in *<ib_install>/examples/security* directory: *readmeta.sql*, *writemeta.sql* and *blindmeta.sql.* These scripts can be run against databases with ISQL to make wholesale changes to a database's system tables access privileges.

- *readmeta.sql* applies the default PUBLIC access privileges: PUBLIC can only select from the system tables, but the database owner, system administrator, and SYSDBA user have full access. This script can be used to return a database that has customized system table privileges back to this default.
- *writemeta.sql* grants all system table privileges to PUBLIC. This is the behavior that existed in InterBase 6.0 and earlier.
- *blindmeta.sql* revokes all system table privileges from PUBLIC. This prevents any PUBLIC user from querying the system tables, including InterBase and third-party utilities run by PUBLIC users. For example, **gstat**, **gbak**, QLI and IBConsole would not be able to query system metadata. This script allows developers to protect their intellectual property by hiding the database design of tables, stored procedures and triggers from the general public and competitors. Blind access makes it difficult, if not impossible, for a general user to generate ad hoc queries against a database.

  A database with blind access does not prevent any user from using InterBase Data Definition Language (DDL) to define new database objects. It just prevents a user from querying or writing to the system tables directly.

  isc_blob_lookup_desc() and isc_array_lookup_bounds() Two client-side APIs, *isc_blob_lookup_desc*() and *isc_array_lookup_bounds*(), cannot execute without SELECT metadata privileges, because the APIs directly query some InterBase system tables. Thus databases that have had *blindmeta.sql* run against them are not able to execute these APIs for any users except the owner, the system administrator, and SYSDBA.

  Older InterBase clients InterBase 6.0 and previous InterBase kits cannot access a database on behalf of a user if that user has no privileges to the system tables. Thus an InterBase developer who runs *blindmeta.sql* on an InterBase database cannot ship that database to customers with InterBase 6.0 or older runtime kits and expect those users to be able to access the database. The developer would have to run *readmeta.sql* against a copy of the database and ship that to customers who have older InterBase runtimes.

Migration issues

The InterBase engine automatically installs the default (SELECT-only) SQL privileges for PUBLIC on the system tables when attaching ODS 10.0 or older databases. Thus if all users must be able to write to database metadata, *writemeta.sql*will have to be run against each database to restore that behavior.

SQL privileges

Connecting to a database does not automatically include privileges to modify or even view data stored within that database. Privileges must be granted explicitly; users cannot access any database objects until they have been granted privileges. Privileges granted to PUBLIC apply to all users.

For full description of syntax of SQL privileges, see entries for GRANT and ROLE in the *Language Reference* and *Data Definition Guide*.

Groups of users

InterBase implements features for assigning SQL privileges to groups of users. SQL roles are implemented on a per-database basis. UNIX groups are implemented on a server-wide basis, using the UNIX *group* mechanism.

SQL roles

InterBase supports SQL group-level security as described in the *ISO-ANSI Working Draft for Database Language*. For syntax of SQL ROLEs, see the *Language Reference* and *Data Definition Guide*.

Implementing roles is a four-step process.

1. Declare the role with CREATE ROLE.

```
CREATE ROLE sales;
```

2. Assign privileges on specific tables and columns to the role using the GRANT statement.

```
GRANT UPDATE ON table1 TO sales;
```

3. Grant the role to users, again with the GRANT statement.

```
GRANT sales TO user1, user2, user3;
```

4. Finally, to acquire the privileges assigned to a role, users must specify the role when connecting to a database.

```
CONNECT 'foo.ib' USER 'user1' PASSWORD 'peanuts' ROLE sales;
```

User1 now has update privileges on TABLE1 for the duration of the connection.

A user can belong to only one role per connection to the database and cannot change role while connected. To change role, the user must disconnect and reconnect, specifying a different role name.

You can adopt a role when connecting to a database by any one of the following means:

- **To specify a role when attaching to a database through IBConsole ISQL**, display the Database Connect dialog and type a rolename in the Role field.
- **To specify a role programmatically upon connection using the InterBase API**, use the dpb parameter *isc_dpb_sql_role_name*. See chapter 4 of the *API Guide*.
- **To specify a role for a connection made by an embedded SQL application or isql session**, use the ROLE *rolename* clause of the CONNECT statement. See the statement reference for CONNECT in the *Language Reference*.

Note Applications using BDE version 5.02 or later, including Delphi, JBuilder, and C++Builder, have a property by which they can specify a role name. Also, the ODBC driver that currently ships with InterBase also recognizes roles.

UNIX groups

Operating system-level groups are implicit in InterBase security on UNIX, similarly to the way UNIX users automatically supplement the users in the InterBase security database. For full description of usage and syntax of using UNIX groups with InterBase security, see the *Language Reference* and *Data Definition Guide*.

Note Integration of UNIX groups with database security is not a SQL standard feature.

Other security measures

InterBase provides some restrictions on the use of InterBase tools in order to increase security. In addition, there are things that you can do to protect your databases from security breaches. This section describes these options.

Restriction on using InterBase tools

As a security measure, InterBase requires that only the owner of a database or SYSDBA can execute **gbak**, **gstat**, and **gfix**.

- Only the database owner or SYSDBA can use **gbak** to back up a database. Anyone can restore a database, because there is no concept of an InterBase user for a backup file. However, only the owner or SYSDBA can restore a database over an existing database. For security purposes, make sure that your backup files are stored in a secure location. This prevents unauthorized persons from restoring databases and gaining access to them.
- On UNIX platforms, there is a further constraint on **gstat**: to run **gstat**, you must have system-level read access to the database file. To access the database with **gstat**, you must either be logged into the account running the InterBase server ("interbase" or "root") or someone must change the permissions on the database file to include read permission for your Group.

Protecting your databases

You can take several steps to increase the security of your databases and other files on your system:

- UNIX and Linux systems: Before starting the InterBase server, log in as user "interbase" (or "interbas", if user names longer than eight characters are not allowed), rather than "root" (only these users can start the server). This restricts the ability of other users to accidentally or intentionally access or overwrite sensitive files such as the password file. Start the InterBase server while you are logged on as user "interbase".
- Windows NT/2000 systems: When the InterBase server is run as a service, you can protect a database against unauthorized access from outside InterBase (such as by a **copy** command), by making the database files readable only by the system account, under which services run. However, if you make the database readable only by the system account, remote access to the database must be by TCP/IP, not by NetBEUI.
- Windows 98/ME systems: When security is an issue, it is inappropriate to run an InterBase server on Windows 98 at all. Due to the lack of a file-permission system, an attacker can readily duplicate the database with the **copy** command or with a point-and-click equivalent. While sometimes a stolen copy of a database will be corrupt and unusable, it might turn out to be adequate, especially if update activity on the database is relatively infrequent, or if the attacker can try copying repeatedly.
- Because anyone can restore a backed up database, it is wise to keep your backup files in a directory with restricted access. (On a Windows 98/ME platform, you can either move backup files to physical media such as tape or high-density removable drives and store these securely, or move the backup files to a system that restricts directory access). On UNIX, only the backup file itself, not the directory in which it resides, needs to have permissions restricted to prevent reading by unauthorized persons.

For example, if all of the following are true:

- the backup file has permission 600 (`rw-------`) or 640 (`rw-r-----`)
- only trusted persons belong to the groups
- the directory has permission `rwxr-xr-x`

then persons other than the responsible owner and group can see that the backup file is there, but they cannot get at it. If the user or backup script issues the command "umask 077" (or 027, as appropriate) before running **gbak**, unauthorized persons will not be able to access the backup file, no matter what the permissions on the directory. (Of course the directory should not be writable by "other"; that would permit other persons to delete the backup file.)

User administration with IBConsole

User administration is accomplished through the User Information dialog where you are able to add, modify, view and delete users. User administration can only be performed after logging in to the server.

Displaying the User Information dialog

You can use any of the following methods to access the User Information dialog:

- Select a logged in server or any branch under the server hierarchy from the list of registered servers in the Tree pane; choose **Server | User Security**.
- Select a logged in server from the list of registered servers in the Tree pane. Double-click User Security in the Work pane or right-click the selected server and choose User Security from the context menu.
- Select Users under the desired server in the Tree pane to display a list of valid users in the Work pane. Double-click a user name to display the User Information dialog.

Figure 5.1 User information dialog

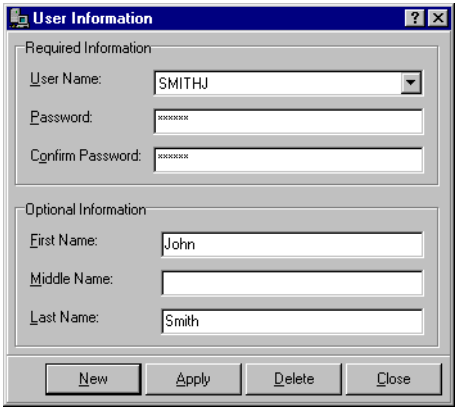

Adding a user

Use the User Information dialog to add new users. To access this dialog follow one of the methods described in "Displaying the User Information dialog" on page 5-9.

To add a new user:

1 Display the User Information dialog in one of the following two ways:

- Select a logged in server or any branch under the server hierarchy from the list of registered servers in the Tree pane; choose **Server | User Security**.

- Select a logged in server from the list of registered servers in the Tree pane. Double-click User Security in the Work pane or right-click the selected server and choose User Security from the context menu.
- Select Users under the desired server in the Tree pane to display a list of valid users in the Work pane. Double-click a user name to display the User Information dialog.

2 Click New. The New and Delete buttons are disabled and the Close button changes to a Cancel button.

3 Type the new username in the User Name text field.

4 Type the user's password in both the Password and the Confirm Password text fields.

5 Add any desired optional information in the corresponding text fields. Each of the optional text fields can be up to 32 characters.

6 Click Apply to add the new user to the security database or click Cancel to abandon your changes.

Note Usernames can be up to 31 characters long and are *not* case sensitive. Passwords *are* case-sensitive and only the first eight characters are significant. InterBase does not allow you to create usernames or passwords containing spaces.

Modifying user configurations

Use the User Information dialog to modify user configurations. To access this dialog follow one of the methods described in "Displaying the User Information dialog" on page 5-9.

To modify a user's details:

1 Display the User Information dialog in one of the following two ways:

- Select a logged in server or any branch under the server hierarchy from the list of registered servers in the Tree pane; choose **Server | User Security** to display the User Information dialog.
- Select a logged in server from the list of registered servers in the Tree pane. Double-click User Security in the Work pane or right-click the selected server and choose User Security from the context menu.
- Select Users under the desired server in the Tree pane to display a list of valid users in the Work pane. Double-click a user name to display the User Information dialog.

2 From the User Name drop down list, select the user whose configuration you wish to modify. The user's details display. You can also type the first letter of the desired username in the User Name drop down list to quickly scroll to usernames beginning with that letter. By repeatedly typing that same letter, you can scroll through all usernames that begin with that letter.

3 Change any of the text fields except the User Name. If you change the password, you must enter the same password in the Password text field and the Confirm Password text field.

4 Click the Apply button to save your changes.

Note You may not modify a username. The only way to change a username is to delete the user and then add a user with the new name.

Deleting a user

Use the User Information dialog to removed users from the security database. To access this dialog follow one of the methods described in "Displaying the User Information dialog" on page 5-9.

To remove a user account:

1 Display the User Information dialog in one of the following two ways:

- Select a logged in server or any branch under the server hierarchy from the list of registered servers in the Tree pane; choose **Server | User Security**.
- Select a logged in server from the list of registered servers in the Tree pane. Double-click User Security in the Work pane or right-click the selected server and choose User Security from the context menu.

2 Select the user you wish to delete from the User Name drop down list. You can also type the first letter of the desired username in the User Name drop down list to quickly scroll to usernames beginning with that letter. By repeatedly typing that same letter, you can scroll through all usernames that begin with that letter.

3 Click Delete. A confirmation dialog inquires, "Do you wish to delete user *username*?" If you choose OK, the user is removed and is no longer authorized to access databases on the current server.

Important Although it is possible for the SYSDBA to delete the SYSDBA user, it is strongly not recommended because it will no longer be possible to add new users or modify existing user configurations If you do delete the SYSDBA user, you must reinstall InterBase to restore the InterBase security database (*admin.ib* by default).

User administration with the InterBase API

The InterBase API includes three functions that permit authors of InterBase applications to add, delete, and modify users programmatically using three API functions: *isc_add_user*( ), *isc_delete_user*( ), and *isc_modifiy_user*( ). These functions are deprecated in InterBase 6 and later, however, because they are replaced by functions in the InterBase Services API.

The InterBase Services API provides a much broader and more robust set of tools for programmatically managing users in the security database. See Chapter 12, "Working with Services" in the *API Guide* for details and examples of using the Services API functions.

For programmers using Delphi and C++ Builder, the IBX components include components for managing users. For more information on using the IBX components, refer to the *Developer's Guide*.

Using gsec to manage security

The InterBase command-line security utility is **gsec**. This utility is used in conjunction with the InterBase security database (*admin.ib* by default) to specify user names and passwords for an InterBase server. This tool duplicates the functionality of **Server | User Security** in IBConsole for Windows.

The security database resides in the InterBase install directory. To connect to a database on the server, users must specify a user name and password, which are verified against information stored in the security database. If a matching row is found, the connection succeeds.

Important Only the SYSDBA can run **gsec**. To do this, use one of the following methods:

- Invoke the command as:

```
gsec -user sysdba -password masterkey
```

- Define the ISC_USER and ISC_PASSWORD environment variables for SYSDBA before you invoke the command.
- Run **gsec** when you are logged in as root on UNIX or Administrator on Windows NT/2000.

To use **gsec** interactively, type **gsec** at the command prompt. The prompt changes to `GSEC>`, indicating that you are in interactive mode. To quit an interactive session, type `QUIT`.

Running gsec remotely

You can use **gsec** on a client host to administer users in a security database on a remote server. Use the **-database** option with a remote database specification to connect to a remote InterBase security database. For example:

```
gsec -database jupiter:/usr/interbase/admin.ib
```

Using gsec commands

The following table summarizes **gsec** commands. The initial part of each command is required. The part in brackets is optional.

Table 5.2 Summary of **gsec** commands

Command	Description
di[splay]	Displays all rows of the InterBase security database (*admin.ib* by default)
di[splay] ***name***	Displays information only for user *name*
a[dd] ***name*** **-pw password** [***option argument***] [***option argument ...***]	Adds user *name* to the security database with password *string*. Each option and corresponding argument specifies other data associated with the user, as shown in Table 5.3, "gsec options"
mo[dify] ***name*** [***options***]	Like **add**, except that *name* already exists in the security database
de[lete] ***name***	Deletes user *name* from the security database
h[elp] or **?**	Displays **gsec** commands and syntax
q[uit]	Quits the interactive session

Displaying the security database

To see the contents of the InterBase security database, enter the DISPLAY command at the GSEC> prompt. All the rows in the security database are displayed:

```
GSEC> display
user nameuid  gidfull name
--------------------------------------------
JOHN    123   345John Doe
JANE    124   345Jane Doe
RICH    125   345Richard Roe
```

Note that passwords are never displayed.

Adding entries to the security database

To add users to the security database, use the **add** command:

```
a[dd] name -pw password [options]
```

followed by a user name, the **-pw** option followed by a password, and any other options, as shown in the following table. The password is case sensitive. None of the other parameters are case sensitive.

For each option, the initial letter or letters are required and optional parts are enclosed in brackets. Each option must be followed by a corresponding argument, a string that specifies the data to be entered into the specified column in the InterBase security database (*admin.ib* by default).

Table 5.3 gsec options

Option	Meaning
-password or **-pa** *string*	Password of user who is performing the change
-user *string*	User who is performing the change
-pw *string*	Target user password
-uid *integer*	Target user ID
-gid *integer*	Group ID for target user
-fname *string*	First Name for target user
-mname *string*	Middle Name for target user
-lname *string*	Last Name for target user

Note The **-pa** switch specifies the root or the SYSDBA account password; **-pw** specifies the password for the user being added or modified.

For example, to add user "jones" and assign the password "welcome", enter:

```
GSEC> add jones -pw welcome
```

Use **display** to verify the entry. An unassigned UID or GID defaults to 0:

```
GSEC> display
user name     uid     gid     full name
--------------------------------------------
JONES         0       0
```

For example, to add authorization for a user named Cindi Brown with user name "cbrown" and password "coffee2go", use the following **gsec** command:

```
GSEC> add cbrown -pw coffee2go -fname cindi -lname brown
```

To verify the new entry, display the contents of the security database:

```
GSEC> display
user name     uid     gid     full name
--------------------------------------------
JONES         0       0
CBROWN        0       0       CINDI  BROWN
```

gsec stores the user name in uppercase regardless of how it is entered.

Modifying the security database

To change existing entries in the security database, use the **modify** command. Supply the user name for the entry to change, followed by the option indicating the items to change and the corresponding values to which to change them.

For example, to set the user ID of user “cbrown” to 8 and change the first name to “Cindy”, enter the following commands:

```
GSEC> modify cbrown -uid 8 -fname cindy
```

To verify the changed line, use **display** followed by the user name:

```
GSEC> display cbrown
user name     uid   gid   full name
------------------------------------------
CBROWN        8     0     CINDY BROWN
```

Note To modify a user name, first delete the entry in the security database, then enter the new user name and re-enter the other information.

Deleting entries from the security database

To delete a user’s entry from the security database, use **delete** and specify the user name:

```
GSEC> delete cbrown
```

You can confirm that the entry has been deleted with the **display** command.

Using gsec from a Windows command prompt

To use **gsec** from the Windows NT or Windows 2000 command prompt, precede each command with **gsec** and prefix each **gsec** command with a hyphen (-). For example, to add user “aladdin” and assign the password, “sesame”, enter the following at the command line:

```
C:> gsec -add aladdin -pw sesame
```

To display the contents of the InterBase security database, enter:

```
C:> gsec -display
```

gsec error messages

Table 5.4 **gsec** security error messages

Error Message	Causes and Suggested Actions to Take
Add record error	The **add** command either specified an existing user, used invalid syntax, or was issued without appropriate privilege to run **gsec**. Change the user name or use **modify** on the existing user.
<string> already specified	During an **add** or **modify**, you specified data for the same column more than once. Retype the command.
Ambiguous switch specified	A command did not uniquely specify a valid operation.
Delete record error	The **delete** command was not allowed. Check that you have appropriate privilege to use **gsec**.
Error in switch specifications	This message accompanies other error messages and indicates that invalid syntax was used. Check other error messages for the cause.
Find/delete record error	Either the **delete** command could not find a specified user, or you do not have appropriate privilege to use **gsec**.
Find/display record error	Either the **display** command could not find a specified user, or you do not have appropriate privilege to use **gsec**.
Find/modify record error	Either the **modify** command could not find a specified user, or you do not have appropriate privilege to use **gsec**.
Incompatible switches specified	Correct the syntax and try again.
Invalid parameter, no switch defined	You specified a value without a preceding argument.
Invalid switch specified	You specified an unrecognized option. Fix it and try again.
Modify record error	Invalid syntax for **modify** command. Fix it and try again. Also check that you have appropriate privilege to run **gsec**.

Table 5.4 gsec security error messages (*continued*)

Error Message	Causes and Suggested Actions to Take
No user name specified	Specify a user name after **add**, **modify**, or **delete**.
Record not found for user: <*string*>	An entry for the specified user could not be found. Use **display** to list all users, then try again.
Unable to open database	The InterBase security database does not exist or cannot be located by the operating system.

Chapter

6

Database Configuration and Maintenance

This chapter describes configuration and maintenance issues for individual databases, including the following topics:

- Database files
- On-disk structure (ODS)
- Read-write and read-only databases
- Creating databases
- Backup file properties
- Shadowing
- Setting database properties
- Sweep interval and automated housekeeping
- Configuring the database cache
- Forced writes vs. buffered writes
- Validation and repair
- Shutting down and restarting databases
- Limbo transactions
- gfix command-line tool

Database files

InterBase database files are in many cases self-contained. All the data and indexes are maintained as data structures within one type of file. The transaction log is also kept within this file.

You can extend the functions available in InterBase database metadata by creating libraries of functions compiled in your language of choice. You can compile functions into a dynamic library (called a DLL on Windows, and a shared library on UNIX) and use them in queries, stored procedures, triggers, views, and so on.

Database file size

InterBase database file size is the product of the number of database pages times the page size. The minimum page size is 1 KB, the default page size is 4KB, and the maximum page size is 8KB. Each page can store records only from a single table. You set the database page size when you create a database by using the PAGE SIZE clause of the CREATE DATABASE statement, or its equivalent in IBConsole. You can change the page size when you restore a database using **gbak** or IBConsole.

Database files cannot exceed 4GB total size (2GB on Windows 98 and Linux). When a database file reaches 4GB, InterBase issues a warning and stops writing data to that file. Before the database reaches this point, you should add one or more secondary files using the ALTER DATABASE statement, or IBConsole, or by using **gbak** to back the database up and then restore it.

Note Using **gbak** is the *only* way to reduce the size of the primary database file. When you restore a database, you can specify multiple files without reference to the original file sizes.

Dynamic file sizing

InterBase dynamically expands the last file in a database as needed until it reaches the 4GB limit (2GB on Windows 98 and Linux). This applies to single-file database as well as to the last file of multifile databases. Specifying a LENGTH for the last or only file in a database has no effect.

External files

InterBase permits external files to be used as *external tables*. These tables are limited in their functionality:

- From a database that is in read-write mode, you can execute only SELECT and INSERT statements on external tables. From a read-only database, you can execute only SELECT statement on external tables.
- You cannot define indexes on external tables; they are outside of the control of the multigenerational architecture.

The default location for external files is *<interbase_home>/ext*. InterBase can always find external files that you place here. If you want to place them elsewhere, you must specify the location in the *ibconfig* configuration file using the EXTERNAL_FILE_DIRECTORY entry.

Important For security reasons, it is extremely important that you not place files with sensitive content in the same directory with external tables.

Migration note: If you are migrating from InterBase 6.x or older to InterBase 7.x or newer, and your database includes external table files, you must either move these files to <*interbase_home*>/*ext* or note their locations in *ibconfig* using the EXTERNAL_FILE_DIRECTORY entry

Temporary files

InterBase dynamically creates files in the temporary file space for scratch space during sorting operations involving large amounts of data. See "Managing temporary files" on page 3-12 for details on temporary file use.

File naming conventions

In earlier versions, InterBase database files were given a file extension of *gdb* by convention. InterBase no longer recommends using *gdb* as the extension for database files, since on some versions of Windows ME and Windows XP, any file that has this extension is automatically backed up by the System Restore facility whenever it is touched. On those two platforms, using the *gdb* extension for database names can result in a significant detriment to performance. Linux and Solaris are not affected. InterBase now recommends using *gdb* as the extension for database names.

InterBase is available on a wide variety of platforms. In most cases users in a heterogeneous networking environment can access their InterBase database files regardless of platform differences between client and server machines if they know the target platform's file naming conventions.

Generally, InterBase fully supports each platform's file naming conventions, including the use of node and path names. InterBase, however, recognizes two categories of file specification in commands and statements that accept more than one file name. The first file specification is called the *primary file specification*. Subsequent file specifications are called *secondary file specifications*. Some commands and statements place restrictions on using node names with secondary file specifications. In syntax statements, file specification is denoted as '*filespec*'

Primary file specifications

InterBase syntax always supports a full file specification, including optional node name and full path, for primary file specifications. For example, the syntax notation for CREATE DATABASE appears as follows:

```
CREATE {DATABASE | SCHEMA} 'filespec'
   [USER 'username' [PASSWORD 'password']]
   [PAGE_SIZE [=] int]
   [LENGTH [=] int [PAGE[S]]]
   [DEFAULT CHARACTER SET charset]
```

In this syntax, the *filespec* that follows CREATE DATABASE supports a node name and path specification, including a platform-specific drive or volume specification.

Secondary file specifications

For InterBase syntax that supports multiple file specification, such as CREATE DATABASE, all file specifications after the first one are *secondary*. Secondary file specifications cannot include a node name, but can specify a full path name.

Multifile databases

InterBase supports databases that span multiple files and multiple file systems. You can add additional files to the database without having to take it off line.

The Database Restore task in IBConsole and in the **gbak** command-line utility permit you to create a multifile database. The only way to alter the file size allocation of an existing database is to back up and restore the database file.

Adding database files

You have the option of specifying the size of secondary files in either of two ways: specify the page on which each secondary file starts, or specify the length in database pages of each file. When you specify the size using the LENGTH keyword, do not specify the length of the final file. InterBase sizes the final file dynamically, as needed.

The following **isql** example adds files using STARTING AT syntax:

```
CONNECT ‘first.ib’;

ALTER DATABASE
   ADD FILE 'second.ib' STARTING AT 50000;
```

Altering database file sizes

You cannot use ALTER DATABASE to split an existing database file. For example, if your existing database is 80,000 pages long and you issue the command above, InterBase starts the new database file at page 80,001. The only way to split an existing database file into smaller files is to back it up and restore it. When you restore a database, you are free to specify secondary file sizes at will, without reference to the original file sizes.

The following **isql** example adds a file using LENGTH syntax. *second.ib* will begin on the page following the final page of *first.ib* and will grow to 50,000 database pages. Then InterBase begins writing to *third.ib* and dynamically increases the size as necessary.

```
CONNECT 'first.ib';
ALTER DATABASE ADD FILE 'second.ib' LENGTH 50000
   ADD FILE 'third.ib';
```

InterBase starts writing data to *third.ib* only after *second.ib* file fills up. In the example above, *second.ib* is 50,000 pages long, and begins following the original file. InterBase will begin filling the *third.ib* file after *second.ib* reaches 50,000 pages. Database pages are 4KB each by default and have a maximum size of 8KB.

There is no guarantee that a given table resides entirely in one file or another. InterBase stores records based on available space within database files. Over time, records from a given table tend to spread over all the files in a multifile database.

Maximum number of files

InterBase allows up to 65,536 database files, including shadow files. Note that your operating system might have a much lower limit on the number of simultaneous open files than the **ibserver** process can have.

Application considerations

A multifile database is not the same thing as multiple single-file databases. The tables are all part of the same database they used to be in, but they can be stored across the multiple files. From your application's standpoint, they're all part of the same database and are accessed exactly the same way they would be in a single-file database.

Your application does not need to know about any files except the first one. Any time your database operations access/write data in the secondary files, the InterBase software takes care of it without requiring any special programming from your application. The application attaches to the database by specifying the path of the first file of the database; applications don't change.

Reorganizing file allocation

You can change the sizes of the files of a multifile database when using **gbak** to restore a database. If you need to move a multi-file database to a different disk or directory, use **gbak** to back up the database, then specify the new locations of all secondary files as you restore the database. See "gbak command-line tool" on page 8-13.

Tip Any database in a production environment should include a definition for at least one secondary file, even if the current size of the database does not warrant a multifile database. Data tends to accumulate without bounds, and some day in the future your database might exceed your file system size, or the operating system's maximum file size. By defining a secondary file, you specify what action InterBase takes when the database grows beyond these limits. This means that the database administrator is freed from monitoring the database as it approaches the file size limit.

Networked file systems

An InterBase database must reside on a disk local to the server software that accesses it. The database file (including any secondary files and shadow files) cannot reside on networked or remote file systems (called mapped drives on Windows and NFS file systems on UNIX). External tables and UDF libraries can reside on networked file systems, but this practice is not recommended because networked file systems can suffer from intermittent availability.

On UNIX, the InterBase software detects that a database file is located on an NFS file system. In this case, it invokes the remote access method to contact an InterBase server process running on the host that exported the file system. If there is no InterBase server software running on that node, any connection to the database fails.

On-disk structure (ODS)

Each release of InterBase has characteristic features in its internal file format. To distinguish between the file formats, InterBase records an on-disk structure (ODS) number in the database file. In general, major ODS versions (those incrementing the number to the left of the decimal point) introduce features that are not backward compatible with earlier ODS versions. The InterBase 7 format is ODS 11. InterBase 7 also supports ODS version 10.x, but new features in ODS 11 that are not recognized by earlier software include the BOOLEAN datatype and 68-byte metadata names.

When you create a new database or restore a backup file in the current version of InterBase, the resulting database file has the current ODS version.

Important To upgrade the ODS of an older database, you must back it up using the backup utility for the version of the existing database and then restore it using the current version of InterBase.

Read-write and read-only databases

InterBase databases have two modes: read-only and read-write. At creation, all databases are both readable and writable: they are in *read-write mode*.

Read-write databases

To function in read-write mode, databases must exist on writable media and the **ibserver** process must have write access to the database file. For databases that are in read-write mode, this is true even when they are used only for reading because the transaction states are kept in an internal inventory data structure within the database file. Therefore any transaction against the database requires the ability to write to the transaction inventory.

Under both Windows NT and UNIX, read-write database files must be writable by the user ID for the **ibserver** process. However, the operating environment or file system can be configured to create files that have limited file privileges by default. If you attempt to attach to a database and get an error of "unavailable database," first check to see if the database file's permissions are such that the user ID of the **ibserver** process does not have write privilege on the database file.

Read-only databases

You can change InterBase databases to *read-only mode*. This provides enhanced security for databases by protecting them from accidental or malicious updates and enables distribution on read-only media such as CDROMs. Databases are always in read-write mode at creation time. This feature is independent of dialect. Any ODS 10 or higher database can be set to read-only mode.

You can use **gbak**, **gfix**, or IBConsole to change a database to read-only mode. (See "Making a database read-only" below.)

Properties of read-only databases

- In read-only mode, databases can be placed on CD-ROMs or in read-only file systems as well as on read-write file systems.
- Attempted INSERT, UPDATE, and DELETE operations on a read-only database generate an error. See the "Error Codes and Messages" chapter of the *Language Reference*.
- No metadata changes are allowed in read-only databases.
- Generators in a read-only database do not increment and are allowed only to return the current value. For example, in a read-only database, the following statement succeeds:

  ```
  SELECT GEN_ID(generator_name, 0) FROM table_name;
  ```

  The following statement fails with the error "attempted update on read-only database."

  ```
  SELECT GEN_ID(generator_name, 1) FROM table_name;
  ```

- External files accessed through a read-only database open in read-only mode, regardless of the file's permissions at the file system level.

Making a database read-only

To change the mode of a database between read-write and read-only, you must be either its owner or SYSDBA and you must have exclusive access to a database.

From within InterBase, you can change a read-write database to read-only mode in any of three ways:

- In IBConsole, select the database, display its properties, and edit the mode. For more information, refer to "Setting database properties" on page 6-18.

- Use **gbak** to back up the database and restore it in read-only mode:

```
gbak -create -mode read_only foo.ibk foo.ib
```

- Use **gfix** to change the mode to read-only:

```
gfix -mode read_only foo.ib
```

Important To set a database to read-only mode from any application that uses BDE, ODBC, or JDBC, use the *isc_action_svc_properties*() function in the InterBase Services API.

Tip To distribute a read-write database on a CD-ROM, back it up and put the *database.ibk* file on the CD-ROM. As part of the installation, restore the database to the user's hard disk.

Read-only with older InterBase versions

- A pre-6 InterBase client can access a read-only database to perform SELECTs. No other operation succeeds.
- If a current InterBase client tries to set a pre-6 database to read-only mode, the server silently ignores the request. There is no way to make older databases read-only. You must upgrade them.

Creating databases

You can create databases on local and remote servers using IBConsole with the Create Database dialog.

You can use any of the following methods to access the Create Database dialog:

- In the Tree pane, select a server or anywhere in the branch under the desired server and choose **Database | Create Database**.
- In the Tree pane, right click the Databases branch under the desired server, and select Create Database from the context menu.

Figure 6.1 Create Database dialog

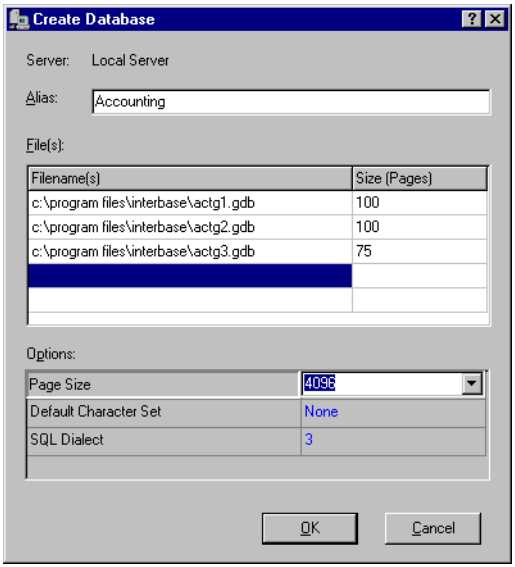

To create a database:

1. Ensure that the server indicated is correct. If it is not, cancel this dialog and re-initiate it under the correct server.
2. Type an Alias name for the new database in the Alias text field.
3. Enter one or more filenames which will make up the database, specifying the number of pages required for each file. To insert a new row into the Files table, move to the last row and column of the table and type Ctrl-Tab.

   When entering a filename, make sure to include the file path unless you wish to default the file to the working directory.

   Note Database files must reside on a local drive.
4. You can specify create options by entering a valid value, by clicking the option value and choosing a new value from a drop down list of values or by double-clicking the option value to rotate its value to the next in the list of values. For more information, see "Database options" below.

   To create a basic database without any options, leave all options blank.
5. Click OK to create the specified database.

Important The alias name that you specify when creating a database references the necessary database file information associated with the database. When performing database configuration and maintenance, you need only specify the alias name, not the actual database filename. If the database spans multiple files, the server uses the header page of each file to locate additional files.

Database options

The database options that you can set are Page Size, Default Character Set, and SQL dialect.

Page size

InterBase supports database page sizes of 1024, 2048, 4096, and 8192 bytes. The default is 4096 bytes.

Default character set

See "Character Set" in Table 9.2 for a detailed explanation of character sets.

For more information about creating databases, see the *Language Reference*. See the *Data Definition Guide* for an explanation of character sets.

SQL dialect

An InterBase database SQL dialect determines how double quotes, large exact numerics, and certain datatypes such as SQL DATE, TIME, and TIMESTAMP are interpreted. In most cases you should create databases in dialect 3 in order to have access to all current InterBase features.

Changing a database dialect from 1 to 3 may require some preparation if it contains DATE datatypes, DECIMAL or NUMERIC datatypes with precision greater than 9, or has strings that are in double quotes rather than single quotes. For more information about dialects, refer to "Understanding SQL dialects" in the migration appendix of the *Operations Guide.*

To change the database dialect:

1 Highlight the database in the Tree pane and perform one of the following actions:

- Choose **Database | Properties**.
- Right-click and choose Properties from the context menu.
- Double-click Properties in the Work pane.

2 Click the General tab and change the SQL dialect in the Options field.

Tip To suppress the display of system tables in IBConsole, deselect System Data from the View menu.

Dropping databases

You can drop databases using IBConsole. Dropping a database deletes the current database and database alias, removing both data and metadata.

To drop a database:

1 Select the database you wish to drop in the Tree pane.

2 Choose **Database | Drop Database** or select Drop Database from the Work pane.

3 A dialog asks you to confirm that you wish to delete the database. Click Yes if you want to drop the selected database, otherwise click No.

Note A database can be dropped only by its creator or SYSDBA.

Important Dropping a database deletes all data and metadata in the database.

Backup file properties

You can view and modify backup file information in IBConsole with the Backup Alias Properties dialog. You can access this dialog with either of the following methods:

- Expand Backup in the Tree pane, select a backup alias, and double-click Modify Backup Alias from the Work pane.
- Right-click a backup alias in the Tree pane and choose Modify Backup Alias from the context menu.

Figure 6.2 Backup alias properties

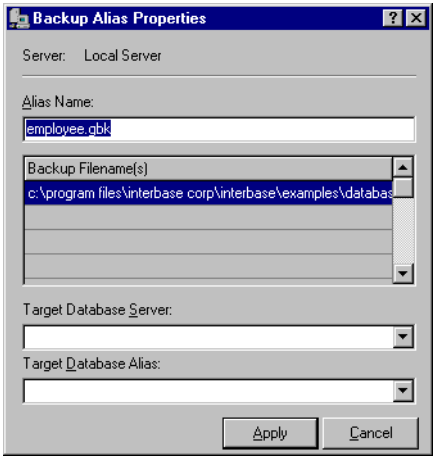

To edit backup file properties:

1 Enter a new backup alias name in the Alias Name text field.

2 Add, remove, or modify the backup filenames and corresponding file sizes associated with the backup in the backup files table. When specifying filenames, be sure to include the file path where the file is located.

To add a new row to the backup files table, move to the last row and column of the table and type Ctrl-Tab. To remove a file from the backup file list, delete (blank out) the values in the table.

3 Select a server from the Target Database Server drop down list. You can also type the server name in the edit portion of the drop down list.

4 Select a database alias from the Target Database Alias drop down list. You can also type the alias name in the edit portion of the drop down list

5 Click Apply to save your changes.

Removing database backup files

You can remove database backup files in IBConsole with either of the following methods:

- Expand Backup in the Tree pane and select a backup alias and double-click Delete Alias from the Work pane.
- Right-click a backup alias in the Tree pane and choose Delete Alias from the context menu.

A dialog asks you to confirm that you wish to remove the selected backup file. Click Yes if you want to delete the backup file, otherwise click No.

Shadowing

InterBase lets you recover a database in case of disk failure, network failure, or accidental deletion of the database. The recovery method is called *disk shadowing*, or sometimes just *shadowing*. This chapter describes how to set up and use shadowing.This section describes the various tasks involved in shadowing, as well as the advantages and limitations of shadowing.

Tasks for shadowing

The main tasks in setting up and maintaining shadowing are as follows:

- Creating a shadow.

Shadowing begins with the creation of a shadow. A shadow is an identical, physical copy of a database. When a shadow is defined for a database, changes to the database are written simultaneously to its shadow. In this way, the shadow always reflects the current state of the database. For information about the different ways to define a shadow, see **"Creating a shadow"** on page 6-14.

- Activating a shadow.

If something happens to make a database unavailable, the shadow can be activated. *Activating* a shadow means it takes over for the database; the shadow becomes accessible to users as the main database. Activating a shadow happens

either automatically or through the intervention of a database administrator, depending on how the shadow was defined. For more information about activating a shadow, see "Activating a shadow" on page 6-17.

- Deleting a shadow.

If shadowing is no longer desired, it can be stopped by deleting the shadow. For more information about deleting a shadow, see "Dropping a shadow" on page 6-17.

- Adding files to a shadow.

A shadow can consist of more than one file. As shadows grow in size, files can be added to accommodate the increased space requirements. For more information about adding shadow files, see "Adding a shadow file" on page 6-18.

Advantages of shadowing

Shadowing offers several advantages:

- Recovery is quick. Activating a shadow makes it available immediately.
- Creating a shadow does not require exclusive access to the database.
- Shadow files use the same amount of disk space as the database. Log files, on the other hand, can grow well beyond the size of the database.
- You can control the allocation of disk space. A shadow can span multiple files on multiple disks.
- Shadowing does not use a separate process. The database process handles writing to the shadow.
- Shadowing can run behind the scenes and needs little or no maintenance.

Limitations of shadowing

Shadowing has the following limitations:

- Shadowing is not an implementation of *replication*. Shadowing is one-way writing, duplicating every write operation on the master database. Client applications cannot access the shadow file directly. For information on replication, see Chapter 12, "Data Replication"
- Shadowing is useful only for recovery from hardware failures or accidental deletion of the database. User errors or software failures that corrupt the database are duplicated in the shadow.
- Recovery to a specific point in time is not possible. When a shadow is activated, it takes over as a duplicate of the database. Shadowing is an "all or nothing" recovery method.

- Shadowing can occur only to a local disk. Shadowing to a NFS file system or mapped drive is not supported. Shadowing to tape or other media is unsupported.

Creating a shadow

A shadow is created with the CREATE SHADOW statement in SQL. Because this does not require exclusive access, it can be done without affecting users. For detailed information about CREATE SHADOW, see the *Language Reference*.

Before creating a shadow, consider the following topics:

- The location of the shadow

A shadow should be created on a different disk from where the main database resides. Because shadowing is intended as a recovery mechanism in case of disk failure, maintaining a database and its shadow on the same disk defeats the purpose of shadowing.

- Distributing the shadow

A shadow can be created as a single disk file called a shadow file or as multiple files called a shadow set. To improve space allocation and disk I/O, each file in a shadow set can be placed on a different disk.

- User access to the database

If a shadow becomes unavailable, InterBase can either deny user access to the database until shadowing is resumed, or allow access even though database changes are not being shadowed. Depending on which database behavior is desired, the database administrator creates a shadow either in auto mode or in manual mode. For more information about these modes, see **"Auto mode and manual mode**" on page 6-16.

- Automatic shadow creation

To ensure that a new shadow is automatically created, create a conditional shadow. For more information, see **"Conditional shadows**," in this chapter.

The next sections describe how to create shadows with various options:

- Single-file or multifile shadows
- Auto or manual shadows
- Conditional shadows

These choices are not mutually exclusive. For example, you can create a single-file, conditional shadow in manual mode.

Creating a single-file shadow

To create a single-file shadow for database *employee.gdb*, enter:

```
SQL> CREATE SHADOW 1 '/usr/interbase/examples/employee.shd';
```

The name of the shadow file is *employee.shd*, and it is identified by the number 1. Verify that the shadow has been created by using the **isql** command SHOW DATABASE:

```
SQL> SHOW DATABASE;
   Database: employee.gdb
    Shadow 1: '/usr/interbase/examples/employee.shd' auto
   PAGE_SIZE 4096
   Number of DB pages allocated = 392
   Sweep interval = 20000
```

The page size of the shadow is the same as that of the database.

Creating a multifile shadow

If your database is large, you can shadow it to a multifile shadow, spreading the shadow files over several disks. To create a multifile shadow, specify the name and size of each file in the shadow set. As with multifile databases, you have the option of specifying the size of secondary files in either of two ways: specify the page on which each secondary file starts, or specify the length in database pages of each file. When you specify the size using the LENGTH keyword, do not specify the length of the final file. InterBase sizes the final file dynamically, as needed.

For example, the following example creates a shadow set consisting of three files. The primary file, *employee.shd*, is 10,000 database pages in length. The second file is 20,000 database pages long, and the final file grows as needed.

```
SQL> CREATE SHADOW 1 'employee.shd' LENGTH 10000
CON> FILE 'emp2.shd' LENGTH 20000
CON> FILE 'emp3.shd';
```

Instead of specifying the page length of secondary files, you can specify their starting page. The following example creates the same shadows as the previous example:

```
SQL> CREATE SHADOW 1 'employee.shd'
CON> FILE 'emp1.shd' STARTING AT 10000
CON> FILE 'emp2.shd' STARTING AT 30000;
```

In either case, you can use SHOW DATABASE to verify the file names, page lengths, and starting pages for the shadow just created:

```
SQL> SHOW DATABASE;
   Database: employee.gdb
    Shadow 1: '/usr/interbase/examples/employee.shd' auto length 10000
    file /usr/interbase/examples/emp1.shd length 2000 starting 10000
    file /usr/interbase/examples/emp2.shd length 2000 starting 30000
   PAGE_SIZE 4096
   Number of DB pages allocated = 392
   Sweep interval = 20000
```

Note The page length you allocate for secondary shadow files need not correspond to the page length of the database's secondary files. As the database grows and its first shadow file becomes full, updates to the database automatically overflow into the next shadow file.

Auto mode and manual mode

A shadow can become unavailable for the same reasons a database becomes unavailable (disk failure, network failure, or accidental deletion). If a shadow becomes unavailable, and it was created in *auto mode*, database operations continue automatically without shadowing. If a shadow becomes unavailable, and it was created in *manual mode*, further access to the database is denied until the database administrator intervenes. The benefits of auto mode and manual mode are compared in the following table:

Table 6.1 Auto vs. manual shadows

Mode	Advantage	Disadvantage
Auto	Database operation is uninterrupted	Creates a temporary period when the database is not shadowed The database administrator might be unaware that the database is operating without a shadow
Manual	Prevents the database from running unintentionally without a shadow	Database operation is halted until the problem is fixed Needs intervention of the database administrator

Auto mode

The AUTO keyword directs the CREATE SHADOW statement to create a shadow in auto mode:

```
SQL> CREATE SHADOW 1 AUTO 'employee.shd';
```

Auto mode is the default, so omitting the AUTO keyword achieves the same result.

In AUTO mode, database operation is uninterrupted even though there is no shadow. To resume shadowing, it might be necessary to create a new shadow. If the original shadow was created as a conditional shadow, a new shadow is automatically created. For more information about conditional shadows, see "Conditional shadows" on page 6-17.

Manual mode

The MANUAL keyword directs the CREATE SHADOW statement to create a shadow in manual mode:

```
SQL> CREATE SHADOW 1 MANUAL 'employee.shd';
```

Manual mode is useful when continuous shadowing is more important than continuous operation of the database. When a manual-mode shadow becomes unavailable, further attachments to the database are prevented. To allow database attachments again, the database owner or SYSDBA must enter the following command:

```
gfix -kill database
```

This command deletes metadata references to the unavailable shadow corresponding to *database*. After deleting the references, a new shadow can be created if shadowing needs to resume.

Conditional shadows

You can define a shadow such that if it replaces a database, the server creates a new shadow file, allowing shadowing to continue uninterrupted. A shadow defined with this behavior is called a *conditional shadow*.

To create a conditional shadow, specify the CONDITIONAL keyword with the CREATE SHADOW statement. For example,

```
SQL> CREATE SHADOW 3 CONDITIONAL 'atlas.shd';
```

Creating a conditional file directs InterBase to automatically create a new shadow. This happens in either of two cases:

- The database or one of its shadow files becomes unavailable.
- The shadow takes over for the database due to hardware failure.

Activating a shadow

When a database becomes unavailable, database operations are resumed by activating the shadow. To do so, log in as SYSDBA or the database owner and use **gfix** with the **-activate** (or **-a**) option.

Important Before activating a shadow, check that the main database is unavailable. If a shadow is activated while the main database is available, the shadow can be corrupted by existing attachments to the main database.

To activate a shadow, specify the path name of its primary file. For example, if database *employee.gdb* has a shadow named *employee.shd*, enter:

```
gfix -a employee.shd
```

After a shadow is activated, you should change its name to the name of your original database. Then, create a new shadow if shadowing needs to continue and if another disk drive is available.

Dropping a shadow

To stop shadowing, use the shadow number as an argument to the DROP SHADOW statement. For example,

```
SQL> DROP SHADOW 1
```

If you need to look up the shadow number, use the **isql** command SHOW DATABASE.

Important DROP SHADOW deletes shadow references from a database's metadata, as well as the physical files on disk. Once the files have been removed from disk, there is no opportunity to recover them. However, a shadow is merely a copy of an existing database, so the new shadow is identical to the dropped shadow.

Adding a shadow file

If a database is expected to increase in size, consider adding files to its shadow. To add a shadow file, first use DROP SHADOW to delete the existing shadow, then use CREATE SHADOW to create a multifile shadow.

The page length you allocate for secondary shadow files need not correspond to the page length of the database's secondary files. As the database grows and its first shadow file becomes full, updates to the database automatically overflow into the next shadow file.

Setting database properties

The Database Properties dialog enables you to display and configure certain database settings. You can access the Database Properties dialog by any of the following methods:

- Select a connected database (or any branch under the database hierarchy) in the Tree pane and choose **Database | Properties**.
- Select a connected database in the Tree pane and double-click Properties in the Work pane.
- Right-click a connected database in the Tree pane and choose Properties from the context menu.

Note If you want to change the database file name, the database must be disconnected before you access the Database Properties dialog. The General tab of this dialog is not available for disconnected databases.

The Database Properties dialog contains two tabs, Alias and General.

Alias tab

The Alias tab of the Database Properties dialog is where you can specify an alias name for a database as well as the file path and file name of the selected database.

Figure 6.3 Database Properties - Alias tab

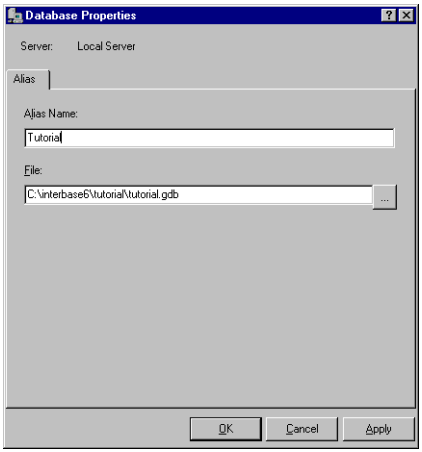

To edit database alias settings:

1. Enter the alias name of the database in the Alias Name text field.
2. Enter database file name, including the path where the file is located, in the File text field. If you prefer, you can also click the browse button ... to locate the file you want.
3. If you need to view or configure the general database settings, click the General tab and see "General tab" below for further information.
4. Once you are finished making changes to the database properties click Apply to save your changes, otherwise click Cancel.

General tab

The General tab of the Database Properties dialog is where you can view such database settings as the database owner, secondary files and their start pages, the number of allocated database pages and the page size. You can also set such options as Forced Writes, Sweep Interval, SQL Dialect and Read Only.

Note If you want to change the database file name, the database must be disconnected before you access the Database Properties dialog. The General tab of this dialog is not available for disconnected databases.

Figure 6.4 Database Properties - General tab

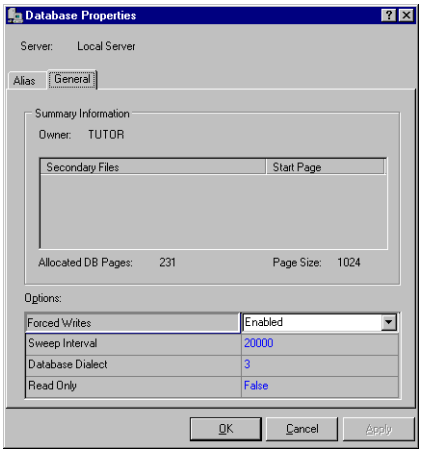

To edit database general options:

1 Choose option values in the Options table. You can specify options by clicking the option value and entering a new value, by choosing a new value from a drop down list of values or by double-clicking the option value to rotate its value to the next in the list of values.

2 If you need to view or configure the database alias settings, click the Alias tab and see "Alias tab" above for further information.

3 Once you are finished making changes to the database properties click Apply to save your changes, otherwise click Cancel.

Table 6.2 General options

Option	Value
Forced Writes	Option values are Enabled and Disabled. See "Forced writes vs. buffered writes" on page 6-26 for further information on forced writes.

Table 6.2 General options

Option	Value
Sweep Interval	The sweep interval is the number of transactions that will occur before an automatic database sweep takes place. You can enter any positive number for the sweep interval, or zero to disable the automatic sweep. See "Sweep interval and automated housekeeping" on page 6-21 for further information on setting the sweep interval.
Database dialect	An InterBase database SQL dialect determines how double quotes, large exact numerics, and certain datatypes such as SQL DATE, TIME, and TIMESTAMP are interpreted. In most cases you should choose dialect 3 in order to have access to all current InterBase features.
Read Only	Option values are True and False. To make the database read only set the Read Only option to True. This prevents users from performing any DML or updates to the database. The default setting for this option is False. See "Making a database read-only" on page 6-7 for more information.

Sweep interval and automated housekeeping

Sweeping a database is a systematic way of removing outdated records. Periodic sweeping prevents a database from growing too large. However, sweeping can also slow system performance.

As a database administrator, you can tune database sweeping, balancing its advantages and disadvantages to best satisfy users' needs.

Overview of sweeping

InterBase uses a multigenerational architecture. This means that multiple versions of data records are stored directly on the data pages. When a record is updated or deleted, InterBase keeps a copy of the old state of the record and creates a new version. This can increase the size of a database.

Garbage collection

To limit the growth of the database, InterBase performs garbage collection by sweeping the database. This process frees up space allocated to outdated record versions. Whenever a transaction accesses a record, outdated versions of that record are collected. Records that were rolled back are not collected. To guarantee that all outdated records are collected, including those that were rolled back, InterBase periodically sweeps the database.

Automatic housekeeping

If a transaction is left in an active (unresolved) state, this is an "interesting" transaction. In a given database's transaction inventory, the first transaction with a state other than committed is known as the Oldest Interesting Transaction (OIT). Automatic housekeeping occurs when the difference between the OIT and the oldest active transaction (OAT) is greater than the sweep interval. By default, this sweep interval is 20,000, but it is configurable (see "Setting the sweep interval" on page 6-22).

Note It is a subtle but important distinction that the automatic sweep does *not* necessarily occur every 20,000 transactions. It is only when the *difference* between the OIT and OAT reaches the threshold. If every transaction to the database is committed promptly, then this difference it is not likely to be great enough to trigger the automatic sweep.

The InterBase server process initiates a special thread to perform this sweep asynchronously, so that the client process can continue functioning, unaffected by the amount of work done by the sweep.

Tip Sweeping a database is not the only way to perform systematic garbage collection. Backing up a database achieves the same result, because the InterBase server must read every record, an action that forces garbage collection throughout the database. As a result, regularly backing up a database can reduce the need to sweep. This enables you to maintain better application performance. For more information about the advantages of backing up and restoring, see "Benefits of backup and restore" on page 8-1.

configuring sweeping

You are able to control several aspects of database sweeping. You can:

- Change the automatic sweep interval.
- Disable automatic sweeping.
- Sweep a database immediately.

The first two functions are performed in the Database Properties dialog. The last is performed with a sweep menu command and is explained in "Performing an immediate database sweep" on page 6-23.

Setting the sweep interval

To set the automatic sweep threshold to *n* transactions:

```
gfix -h n
```

Sweeping a database can affect transaction start-up if rolled back transactions exist in the database. As the time since the last sweep increases, the time for transaction start-up can also increase. Lowering the sweep interval can help reduce the time for transaction start-up.

On the other hand, frequent database sweeps can reduce application performance. Raising the sweep interval could help improve overall performance. The database administrator should weigh the issues for the affected applications and decide whether the sweep interval provides the desired database performance.

To set the sweep interval with IBConsole, refer to "Setting database properties" on page 6-18.

Tip Unless the database contains many rolled back transactions, changing the sweep interval has little effect on database size. As a result, it is more common for a database administrator to tune the database by disabling sweeping and performing it at specific times. These activities are described in the next two sections.

Disabling automatic sweeping

To disable automatic sweeping, set the sweep threshold to zero (0). Disabling automatic sweeping is useful if:

- Maximum throughput is important. Transactions are never delayed by sweeping.
- You want to schedule sweeping at specific times. You can manually sweep the database at any time. It is common to schedule sweeps at a time of least activity on the database server, to avoid competing for resources with clients.

To disable automatic sweeping with IBConsole, refer to "Setting database properties" on page 6-18.

Performing an immediate database sweep

You can perform an immediate database sweep with any of the following methods:

- Right click a connected database in the Tree pane and choose **Maintenance | Sweep** from the context menu
- Select a connected database in the Tree pane and double-click Sweep in the Work pane
- enter the following command:

```
gfix -s
```

This operation runs an immediate sweep of the database, releasing space held by records that were rolled back and by out-of-date record versions. Sweeps are also done automatically at a specified interval.

Sweeping a database does not strictly require it to be shut down. You can perform sweeping at any time, but it can impact system performance and should be done when it inconveniences users the least.

If a sweep is performed as an exclusive operation on the database, there is additional tuning that the procedure performs. As long as there are no outstanding active transactions, the sweep updates the state of data records and the state of the inventory of past transactions. Non-committed transactions are finally rendered obsolete, and internal data structures need not track them in order to maintain snapshots of database versions. The benefit of this is a reduction of memory use, and a noticeable performance improvement.

Configuring the database cache

The *database cache* consists of all database pages (also called buffers) held in memory at one time. *Database cache size* is the number of database pages. You can set the default size of the database cache at three levels:

- server level: applies to all databases
- database level: applies only to a single database (using **gfix** to set the size for a specific database)
- connection level: applies only to a specific **isql** connection

We recommend setting cache size at the database level rather than at the server level. This reduces the likelihood of inappropriate database cache sizes.

In a SuperServer installation (Windows, Linux, or UNIX), every database on a server requires RAM equal to the cache size (number of database pages) times the page size. By default, the cache size is 2048 pages per database and the page size is 4KB. Thus, a single database running at the default setting requires 8MB of memory, but three such databases require 24MB of memory.

In Classic installations, the amount of memory required by a database depends on the number of client attachments. Each client is allotted 75 cache pages, so memory usage is calculated by:

```
cache size × number of attachments × 75
```

Default cache size per database

The **buffers** parameter of the **gfix** utility sets the default number of cache pages for a specific database:

```
gfix -buffers n database_name
```

This sets the number of cache pages for the specified database to *n*, overriding the server value, which by default is 2048 pages.

To run **gfix**, you must be either SYSDBA or the owner of the database.

Default cache size per ISQL connection

To configure the number of cache pages for the duration of one **isql** connection, invoke **isql** with the following option:

```
isql -c n database_name
```

n is the number of cache pages to be used as the default for the session; *n* overrides any values set by DATABASE_CACHE_PAGES or **gfix** and must be greater than 9.

A CONNECT statement entered in an **isql** query accepts the argument CACHE n. (Refer to the discussion of CONNECT in the *Language Reference* manual for a full description of the CONNECT function). For example:

```
ISQL> CONNECT database_name CACHE n;
```

The value *n* can be any positive integer number of database pages. If a database cache already exists in the server because of another attachment to the database, the cache size is increased only if *n* is greater than current cache size.

Setting cache size in applications

InterBase API: use the *isc_dpb_num_buffers* parameter to set cache size in a database parameter buffer (DPB).

IBX: use the *num_buffers* parameter to set cache size in the TIBDatabase's parameter list. For example: *num_buffers*=250. For the parameter to be parsed correctly, there must be no spaces around the = sign.

Default cache size per server

For SuperServer installations, you can configure the default number of pages used for the database caches. By default, the database cache size is 2048 pages per database. You can modify this default by changing the value of DATABASE_CACHE_PAGES in the *ibconfig* configuration file. When you change this setting, it applies to every active database on the server.

You can also set the default cache size for each database using the **gfix** utility. This approach permits greater flexibility, and reduces the risk that memory is overused, or that database caches are too small. *We strongly recommend that you use* **gfix** *to set cache size rather than* DATABASE_CACHE_PAGES.

Verifying cache size

To verify the size of the database cache currently in use, execute the following commands in **isql**:

```
ISQL> CONNECT database_name;
ISQL> SET STATS ON;
ISQL> COMMIT;
```

```
Current memory = 415768
Delta memory = -2048
Max memory = 419840
Elapsed time = 0.03 sec
Buffers = 2048
Reads = 0
Writes 2
Fetches = 2
ISQL> QUIT;
```

The empty COMMIT command prompts **isql** to display information about memory and buffer usage. The "Buffers" line specifies the size of the cache for that database.

Forced writes vs. buffered writes

When an InterBase Server performs forced writes (also referred to as synchronous writes), it physically writes data to disk whenever the database performs an (internal) write operation.

If forced writes are not enabled, then even though InterBase performs a write, the data may not be physically written to disk, since operating systems buffer disk writes. If there is a system failure before the data is written to disk, then information can be lost.

Performing forced writes ensures data integrity and safety, but slow performance. In particular, operations that involve data modification are slower.

Forced writes are enabled or disabled in the Database Properties dialog. For more information, refer to "Setting database properties" on page 6-18.

Validation and repair

In day-to-day operation, a database is sometimes subjected to events that pose minor problems to database structures. These events include:

- Abnormal termination of a database application. This does not affect the integrity of the database. When an application is canceled, committed data is preserved, and uncommitted changes are rolled back. If InterBase has already assigned a data page for the uncommitted changes, the page might be considered an orphan page. Orphan pages are unassigned disk space that should be returned to free space.
- Write errors in the operating system or hardware. These usually create a problem with database integrity. Write errors can cause data structures such as database pages and indexes to become broken or lost. These corrupt data structures can make committed data unrecoverable.

Validating a database

You should validate a database:

- Whenever a database backup is unsuccessful.
- Whenever an application receives a “corrupt database” error.
- Periodically, to monitor for corrupt data structures or misallocated space.
- Any time you suspect data corruption.

Database validation requires exclusive access to the database. Shut down a database to acquire exclusive access. If you do not have exclusive access to the database, you get the error message:

```
OBJECT database_name IS IN USE
```

To shut down a database, refer to the directions in “Shutting down a database” on page 6-31.

Validating a database using gfix

To validate a database using gfix, follow these steps:

1. Enter the following command:

```
gfix -v
```

2. If you suspect you have a corrupt database, make a copy of your database using an OS command (**gbak** will not back up corrupt data).
3. Use the **gfix** command to mark corrupt structures in the copied database:

```
gfix -m
```

4. If gfix reports any checksum errors, validate and repair the database again, ignoring any checksum errors:

```
gfix -v -i
```

It may be necessary to validate a database multiple times to correct all the errors.

Validating a database using IBConsole

To validate a database using IBConsole, access the Database Validation dialog by any of the following methods:

- Select a disconnected database in the Tree pane and double-click Validation in the Work pane.
- Right-click a disconnected database in the Tree pane and choose Validation from the context menu.
- Select **Database | Maintenance | Validation**.

Figure 6.5 Database Validation dialog

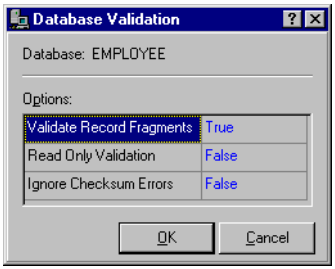

To validate database:

1. Check that the database indicated is correct. If it is not, cancel this dialog and re-initiate the Database Validation dialog under the correct database.
2. Specify which validation options you want by clicking in the right column and choosing True or False from the dropdown list. See Table 6.3, "Validation options" for a description of each option.
3. Click OK if you want to proceed with the validation, otherwise click Cancel.

When IBConsole validates a database it verifies the integrity of data structures. Specifically, it does the following:

- Reports corrupt data structures
- Reports misallocated data pages
- Returns orphan pages to free space

Table 6.3 Validation options

Option	Value
Validate Record Fragments	Option values are True and False. By default, database validation reports and releases only page structures. If the Validate Record Fragments option is set to True, validation reports and releases record structures as well as page structures.
Read Only Validation	Option values are True and False. By default, validating a database updates it, if necessary. To prevent updating, set the Read Only Validation option to True.
Ignore Checksum Errors	Option values are True and False. A checksum is a page-by-page analysis of data to verify its integrity. A bad checksum means that a database page has been randomly overwritten (for example, due to a system crash). Checksum errors indicate data corruption. To repair a database that reports checksum errors, set the Ignore Checksum Errors option to True. This enables IBConsole to ignore checksums when validating a database. Ignoring checksums allows successful validation of a corrupt database, but the affected data may be lost.

Important Even if you can restore a mended database that reported checksum errors, the extent of data loss may be difficult to determine. If this is a concern, you may want to locate an earlier backup copy and restore the database from it.

Repairing a corrupt database

If a database contains errors, they are displayed in the following dialog:

Figure 6.6 Validation report dialog

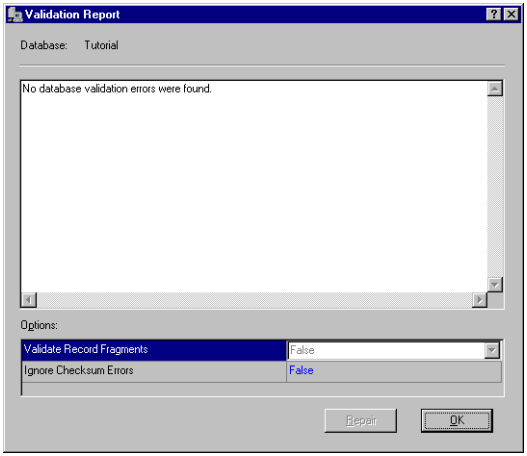

The errors encountered are summarized in the text display area. The repair options you selected in the Database Validation dialog are selected in this dialog also.

To repair the database, choose Repair. This fixes problems that cause records to be corrupt and mark corrupt structures. In subsequent operations (such as backing up), InterBase ignores the marked records.

Note Some corruptions are too serious for IBConsole to correct. These include corruptions to certain strategic structures, such as space allocation pages. In addition, IBConsole cannot fix certain checksum errors that are random by nature and not specifically associated with InterBase.

If you suspect you have a corrupt database, perform the following steps:

1. Make a copy of the database using an operating-system command. Do not use the IBConsole Backup utility or the **gbak** command, because they cannot back up a database containing corrupt data. If IBConsole reports any checksum errors, validate and repair the database again, setting the Ignore Checksum Error option to True.
2. It may be necessary to validate a database multiple times to correct all the errors. Validate the database again, with the Read Only Validation option set to True.

Note Free pages are no longer reported, and broken records are marked as damaged. Any records marked during repair are ignored when the database is backed up.

3. Back up the mended database with IBConsole or **gbak**. At this point, any damaged records are lost, since they were not included during the back up. For more information about database backup, see **Chapter 8, "Database Backup and Restore."**
4. Restore the database to rebuild indexes and other database structures. The restored database should now be free of corruption.

Verify that restoring the database fixed the problem by validating the restored database with the Read Only Validation option set to True.

Shutting down and restarting databases

Maintaining a database often involves shutting it down. Only the SYSDBA or the owner of a database (the user who created it) can shut it down. The user who shuts down the database then has exclusive access to the database.

Exclusive access to a database is required to:

- Validate and repair the database.
- Add or drop a foreign key on a table in the database.
- Add a secondary database file.

After a database is shut down, the database owner and SYSDBA are still able to connect to it, but any other user attempting to connect gets an error message stating that the database is shut down.

Shutting down a database

To shut down a database, select a connected database from the Tree pane and double-click Shutdown in the Work pane or choose **Database | Maintenance | Shutdown** to display the Database Shutdown dialog:

Figure 6.7 Database shutdown dialog

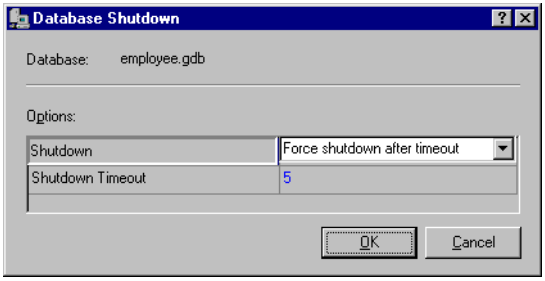

Shutdown timeout options

You can specify a timeout value by selecting a new value from the drop down list of values or by typing the value in the edit portion of the drop down list. Timeout values can range from 1 minute to 500 minutes.

Shutdown options

You can specify shutdown options by selecting a new value from the drop down list of values. Shutdown option values include: Deny New Connections While Waiting, Deny New Transactions While Waiting, and Force Shutdown After Timeout.

Deny new connections while waiting

This option allows all existing database connections to complete their operations unaffected. IBConsole shuts down the database after all processes disconnect from the database. At the end of the timeout period, if there are still active connections, then the database is not shut down.

This prevents any new processes from connecting to the database during the timeout period. This enables current users to complete their work, while preventing others from beginning new work.

Suppose the SYSDBA needs to shut down database *orders.ib* at the end of the day (five hours from now) to perform routine maintenance. The Marketing department is currently using the database to generate important sales reports.

In this case, the SYSDBA would shut down *orders.ib* with the following parameters:

- Deny New Connections.
- Timeout of 300 minutes (five hours).

These parameters specify to deny any new database connections and to shut down the database any time during the next five hours when there are no more active connections.

Any users who are already connected to the database are able to finish processing their sales reports, but new connections are denied. During the timeout period, the SYSDBA sends out periodic broadcast messages asking users to finish their work by 6 p.m.

When all users have disconnected, the database is shut down. If all users have not disconnected after five hours, then the database is not shut down. Because the shutdown is not critical, it is not forced.

It would be inappropriate to deny new transactions, since generating a report could require several transactions, and a user might be disconnected from the database before completing all necessary transactions. It would also be inappropriate to force shutdown, since it might cause users to lose work.

Deny new transactions while waiting

This option allows existing transactions to run to normal completion. Once transaction processing is complete, IBConsole shuts down the database. Denying new transactions also denies new database connections. At the end of the timeout period, if there are still active transactions, then the database is not shut down.

This is the most restrictive shutdown option, since it prevents any new transactions from starting against the database. This option also prevents any new connections to the database.

Suppose the SYSDBA needs to perform critical operations that require shutdown of the database *orders.ib*. This is a database used by dozens of customer service representatives throughout the day to enter new orders and query existing orders.

At 5 p.m., the SYSDBA initiates a database shutdown of *orders.ib* with the following parameters:

- Deny New Transactions.
- Timeout of 60 minutes.

These parameters deny new transactions for the next hour. During that time, users can complete their current transactions before losing access to the database. Simply denying new connections would not be sufficient, since the shutdown cannot afford to wait for users to disconnect from the database.

During this hour, the SYSDBA sends out periodic broadcast messages warning users that shutdown is happening at 6 p.m and instructs them to complete their work. When all transactions have been completed, the database is shut down.

After an hour, if there are still any active transactions, IBConsole cancels the shutdown. Since the SYSDBA needs to perform database maintenance, and has sent out numerous warnings that a shutdown is about to occur, there is no choice but to force a shutdown.

Force shutdown after timeout

With this option, there are no restrictions on database transactions or connections. As soon as there are no processes or connections to the database, IBConsole shuts down the database. At the end of the timeout period, if there are still active connections, IBConsole rolls back any uncommitted transactions, disconnects any users, and shuts down the database.

If critical database maintenance requires a database to be shut down while there are still active transactions, the SYSDBA can force shut down. This step should be taken only if broadcast messages have been sent out to users that shutdown is about to occur. If users have not heeded repeated warnings and remain active, then their work is rolled back.

This option does not deny new transactions or connections during the timeout period. If, at any time during the timeout period, there are no connections to the database, IBConsole shuts down the database.

Important Forcing database shutdown interferes with normal database operations, and should only be used after users have been given appropriate broadcast notification well in advance.

Restarting a database

After a database is shut down, it must be restarted (brought back online) before users can access it.

To restart a database, select a previously shut down database from the Tree pane and choose **Database | Maintenance | Database Restart** or double-click Database Restart in the Work pane. The currently selected database is brought back online immediately.

Limbo transactions

When committing a transaction that spans multiple databases, InterBase automatically performs a two-phase commit. A *two-phase commit* guarantees that the transaction updates either all of the databases involved or none of them—data is never partially updated.

Note The Borland Database Engine (BDE), as of version 4.5, does not exercise the two-phase commit or distributed transactions capabilities of InterBase, therefore applications using the BDE never create limbo transactions.

In the first phase of a two-phase commit, InterBase prepares each database for the commit by writing the changes from each *subtransaction* to the database. A subtransaction is the part of a multi-database transaction that involves only one database. In the second phase, InterBase marks each subtransaction as committed in the order that it was prepared.

If a two-phase commit fails during the second phase, some subtransactions are committed and others are not. A two-phase commit can fail if a network interruption or disk crash makes one or more databases unavailable. Failure of a two-phase commit causes limbo transactions, transactions that the server does not know whether to commit or roll back.

It is possible that some records in a database are inaccessible due to their association with a transaction that is in a limbo state. To correct this, you must recover the transaction using IBConsole. *Recovering* a limbo transaction means committing it or rolling it back. Use **gfix** to recover transactions.

Recovering transactions

You can recover transactions by any of the following methods:

- Select a connected database in the Tree pane and double-click Transaction Recovery in the Work pane or choose **Database | Maintenance | Transaction Recovery**.
- Right-click a connected database in the Tree pane and choose **Maintenance | Transaction Recovery** from the context menu.

The Transaction Recovery dialog contains two tabs, Transactions and Details. The Transactions tab displays a list of limbo transactions that can then be operated upon to recover - that is, to commit or roll back. You can also seek suggested recovery actions and set current actions to perform on the selected limbo transactions. The Details tab displays detailed information about a selected transaction.

Transaction tab

All the pending transactions in the database are listed in the text area of the Transactions tab. You can rollback, commit, or perform a two-phase commit on such transactions.

Figure 6.8 Transaction Recovery - limbo transactions

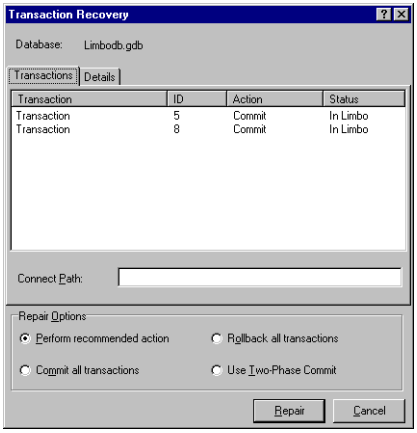

To recover limbo transactions:

1 Select a limbo transaction in the table.

2 The Connect Path text field displays the current path of the database file for the selected transaction, if it is a multi-database transaction. You can change the target database path, if necessary, by overwriting the current path.

The information on the path to the database was stored when the client application attempted the commit. It is possible that the path and network protocol from that machine does not work from the client which is now running IBConsole. Before attempting to roll back or commit any transaction, confirm the path of all involved databases is correct.

When entering the current path, be sure to include the server name and separator indicating communication protocol. To use TCP/IP, separate the server and directory path with a colon (:). To use NetBEUI, precede the server name with either a double backslash (\\) or a double slash (//), and then separate the server name and directory path with either a backslash or a slash.

3 If you want to continue with the transaction recovery process select a repair option and click Repair, otherwise click Cancel. To determine the recommended action, click on the transaction and select the Details tab. For further information about transaction recovery suggestions, see "Details tab" below.

Details tab

The Details tab displays the host server, the remote server, database path, and recommended action: either commit or rollback. If you want to continue with the transaction recovery process select a repair option and click Repair, otherwise click Cancel.

Figure 6.9 Transaction recovery: Details

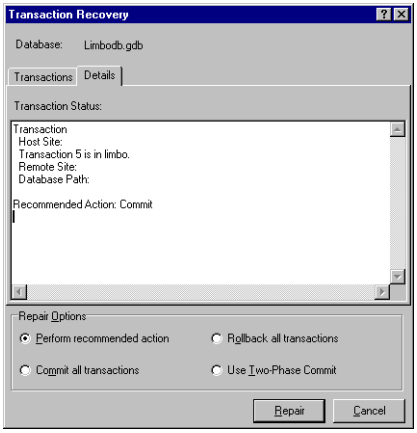

Viewing the administration log

IBConsole displays the administration log file in a standard text display window, the Administration Log dialog, which can be accessed by any of the following methods:

- Select a server (or any branch under the server hierarchy) in the Tree pane and choose **Server | View Logfile**.
- Right-click the desired server in the Tree pane and choose View Logfile from the context menu.
- Under the desired server, select Server Log in the Tree pane and then double-click View Logfile in the Work pane.

Figure 6.10 Administration Log dialog

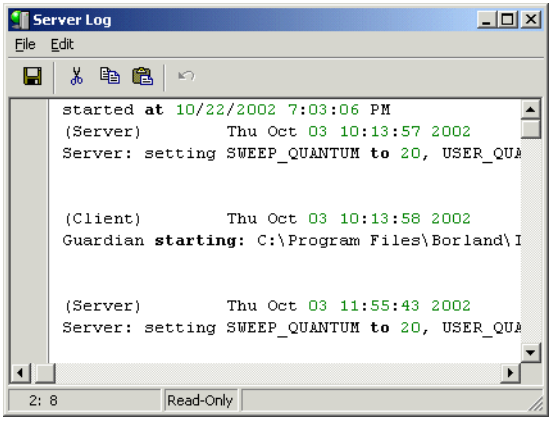

The standard text display window enables you to search for specific text, save the text to a file, and print the text. For an explanation of how to use the standard text display window, see "Standard text display window" on page 2-7.

gfix command-line tool

The **gfix** tool performs a number of maintenance activities on a database, including the following:

- Database shutdown
- Changing database mode to read-only or read-write
- Changing the dialect of a database
- Setting cache size at the database level
- Committing limbo transactions
- Mending databases and making minor data repairs
- Sweeping databases
- Displaying, committing, or recovering limbo transactions

To run **gfix**, you must attach as either SYSDBA or the owner of the database. Most of these actions can also be performed through IBConsole.

Syntax `gfix [options] db_name`

Options In the OPTION column of the following table, only the characters outside the brackets ([ ]) are required. You can specify additional characters up to and including the full option name. To help identify options that perform similar functions, the TASK column indicates the type of activity associated with an option.

Table 6.4 gfix options

Option	Task	Description
-ac[tivate]	Activate shadows	Activate shadows when the database dies. NOTE: syntax is gfix -ac (no database name)
-at[tach] *n*	Shutdown	Used with **-shut** to prevent new database connections during timeout period of *n* seconds; shutdown is canceled if there are still processes connected after *n* seconds
-b[uffers] *n*	Cache buffers	Sets default cache buffers for the database to *n* pages
-ca[che] *n*		Reserved for future implementation
-c[ommit] {ID | all}	Transaction recovery	Commits limbo transaction specified by **ID** or commit all limbo transactions

Table 6.4 **gfix** options (*continued*)

Option	Task	Description
-force *n*	Shutdown	Used with **-shut** to force shutdown of a database after *n* seconds; this is a drastic solution that should be used with caution
-f[ull]	Data repair	Used with **-v** to check record and page structures, releasing unassigned record fragments
-h[ousekeeping] *n*	Sweeping	Changes automatic sweep threshold to *n* transactions • Setting *n* to 0 disables sweeping • Default threshold is 20,000 transactions (see "Overview of sweeping" on page 6-21) • Exclusive access not needed
-i[gnore]	Data repair	Ignores checksum errors when validating or sweeping
-k[ill]	Drop shadows	Drops unavailable shadows. NOTE: syntax is **gfix -k** (no database name)
-l[ist]	Transaction recovery	Displays IDs of each limbo transaction and indicates what would occur if **-t** were used for automated two-phase recovery
-m[end]	Data repair	Marks corrupt records as unavailable, so they are skipped (for example, during a subsequent backup)
-mo[de] [read_write | read_only}	Set access mode	• Sets mode of database to either read-only or read-write • Default table mode is read_write • Requires exclusive access to the database
-n[o_update]	Data repair	Used with **-v** to validate corrupt or misallocated structures; structures are reported but not fixed
-o[nline]	Shutdown	Cancels a **-shut** operation that is scheduled to take effect or rescinds a shutdown that is currently in effect
-pa[ssword] *text*	Remote access	Checks for password *text* before accessing a database
-p[rompt]	Transaction recovery	Used with **-l** to prompt for action during transaction recovery
-r[ollback] {ID | all}	Transaction recovery	Rolls back limbo transaction specified by **ID** or roll back all limbo transactions
-s[weep]	Sweeping	Forces an immediate sweep of the database • Useful if automatic sweeping is disabled • Exclusive access is not necessary
-s[ql_dialect] *n*	Database dialect	Changes database dialect to *n* • Dialect 1 = InterBase 5.*x* compatibility • Dialect 3 = Current InterBase with SQL92 features

Table 6.4 **gfix** options (*continued*)

Option	Task	Description
-sh[ut]	Shutdown	• Shuts down the database • Must be used in conjunction with -**attach**, -**force**, or -**tran**
-t[wo_phase] {ID | all}	Transaction recovery	Performs automated two-phase recovery, either for a limbo transaction specified by **ID** or for all limbo transactions
-tr[an] *n*	Shutdown	Used with -**shut** to prevent new transactions from starting during timeout period of *n* seconds; cancels shutdown if there are still active transactions after *n* seconds
-user *name*	Remote access	Checks for user *name* before accessing a remote database
-v[alidate]	Data repair	Locates and releases pages that are allocated but unassigned to any data structures; also reports corrupt structures
-w[rite] {sync | async}	Database writes	Enables or disables forced (synchronous) writes **sync** enables forced writes; **async** enables buffered writes
-z		Shows version of **gfix** and of the InterBase engine

Examples The following example changes the dialect of the *customer.ib* database to 3:

```
gfix -sql 3 customer.ib
```

The following example changes the customer.ib database to read-only mode:

```
gfix -mo read-only customer.ib
```

gfix error messages

Table 6.5 **gfix** database maintenance error messages

Error Message	Causes and Suggested Actions to Take
Database file name <*string*> already given	A command-line option was interpreted as a database file because the option was not preceded by a hyphen (-) or slash (/). Correct the syntax.
Invalid switch	A command-line option was not recognized.
Incompatible switch combinations	You specified at least two options that do not work together, or you specified an option that has no meaning without another option (for example, **-full** by itself).

Table 6.5 gfix database maintenance error messages (*continued*)

Error Message	Causes and Suggested Actions to Take
More limbo transactions than fit. Try again.	The database contains more limbo transactions than **gfix** can print in a single session. Commit or roll back some of the limbo transactions, then try again.
Numeric value required	The **-housekeeping** option requires a single, non-negative argument specifying number of transactions per sweep.
Please retry, specifying <*string*>	Both a file name and at least one option must be specified.
Transaction number or “all” required	You specified **-commit**, **-rollback**, or **-two_phase** without supplying the required argument.
-mode read_only or read_write	The **-mode** option takes either read_only or read_write as an option.
“read_only” or “read_write” required	The **-mode** option must be accompanied by one of these two arguments.

Chapter

7

Licensing

This chapter describes licensing options and mechanisms for InterBase.

Software activation certificates

To install and use the InterBase software, you must have one or more software activation certificates (also sometimes called "license certificates"). Each certificate bears a unique ID and key pair that enables a specific functionality. During installation, you are required to enter exactly one valid certificate ID and key. Follow the instructions below to enter additional certificate ID/key pairs—and therefore functionality—to InterBase.

The InterBase Media Kit contains the CD-ROM for your platform. You order software activation certificates as needed to enable InterBase functionality. You can purchase additional functionality—such as a greater number of concurrent users—at any time through your sales representative.

Your initial InterBase kit includes licensing for the InterBase client at no additional cost.

License registration tools

On Windows platforms, you can add and remove certificate ID and key numbers using IBConsole.

On all platforms, you can add and remove certificate ID and key numbers using the *iblicense* utility at the system prompt.

Using IBConsole

IBConsole permits you to add, remove, and view certificate ID and key numbers on Windows platforms.

Viewing existing licenses

1. In IBConsole, connect to the server for which you wish to view ID/keys.
2. Expand the node for that server and click Certificates.

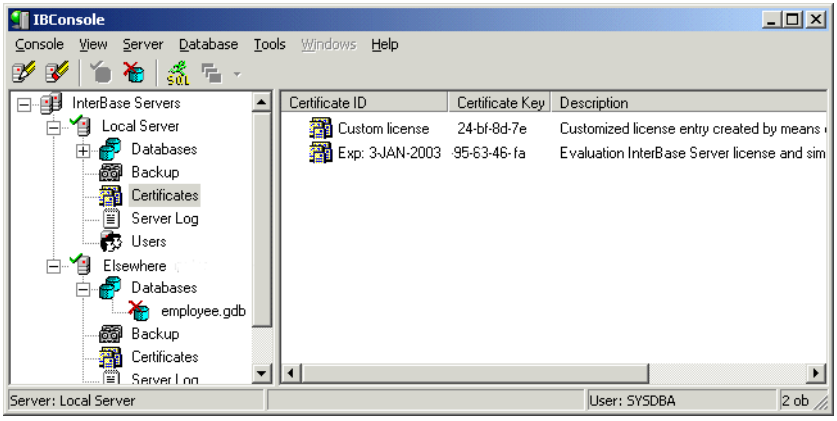

The Work pane displays all current certificate IDs and keys for that server.

Adding a certificate ID/key

1. Have the certificate IDs and keys handy.
2. In IBConsole, connect to the server to which you are adding the certificate ID/keys.
3. Choose **Server** | **Add Certificate** or right-click and choose Add Certificate from the context menu.

Figure 7.1 The Add Certificate dialog

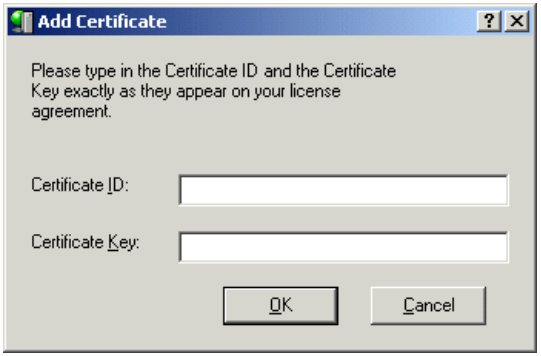

4 Enter the Certificate ID and Certificate Key numbers in the appropriate fields and click OK.

   InterBase adds the license information to the *ib_license.dat* file.

5 Restart the server process for the change to take effect.

Deleting a certificate ID/key

1 In IBConsole, connect to the server that you are deleting the certificate ID/keys from.

2 Expand the node for that server and click Certificates.

3 In the Work pane, highlight the certificate you want to remove.

4 Choose **Server | Remove Certificate**. When you are asked if you're sure you want to remove the certificate, choose Yes.

5 Restart the server process for the change to take effect.

Command-line registration utility

The **iblicense** utility is available on all platforms. When you use it to add certificate ID/key pairs, be sure to have the them handy before you begin.

You must run **iblicense** on the machine where the license file resides.

Command-line syntax

To add only one ID/key, use the following syntax:

```
iblicense -command [-option parameter ...]
```

For example:

```
iblicense -add -key KEYSTRING -id IDNUMBER
```

To add several ID/keys in one session, invoke *iblicense* with no arguments.

```
iblicense Enter
IB_LICEN> command [-option parameter …]
IB_LICEN> command [-option parameter …]
…
IB_LICEN> Quit
```

Important After an add or remove command succeeds, you must restart the server to ensure that the new license is being used.

Removing keys When you remove a key, InterBase saves the current state of the license file in a backup file, named *ib_bckup.dat*.

Options Commands are:

Table 7.1 iblicense commands and their options

Command/option	Description
add -key *keynumber* **-id** *idnumber*	Adds a new license key; the key number and ID number are on your Software Activation Certificate
display	Displays all current keys
remove -key *key*	Removes a key
help	Prints this list
quit	Exits prompt mode
z	Displays the version

Arguments to the **add** and **remove** commands listed above are:

Table 7.2 iblicense options

Option	Description
key	License key value
id	License id value

You can abbreviate a command or option by issuing the initial letters until you have uniquely identified it.

Once the command has been issued, the utility reports whether or not the operation succeeded.

Examples The following are examples of using the *iblicense* utility.

- Adding a key:

```
iblicense -add -key 00-b-90-21 -id 50-60--71-XXXx-01234
```

- Removing a key:

```
iblicense -r -k 00-b-90-21
```

- Displaying keys:

```
iblicense -display
```

- Displaying help:

```
iblicense -h
```

Note The example ID/key pairs above are fictional and do not enable a server.

Available certificates

Each software activation certificate enables a specific piece of functionality, which is identified in the box that contains the Certificate ID and key numbers.

Table 7.3 Certificate keys available for InterBase

Certificate	Functionality
Server Activation	Activates local access to the server for one (1) user
Metadata	Allows database metadata manipulation for the server; that is, SQL statements CREATE, ALTER, and DROP
Remote Access	Enables the server to accept requests for database access from remote clients
Client Capability	Allows the software to act as a client and to connect to other InterBase servers
Desktop InterBase	Activates local access to the server for two users, with metadata and client capabilities
Simultaneous Users	Regulates the number of clients that can connect to this server at the same time; clients can be local or remote
Internet Access	Allows use of the InterBase server with Web servers
Per-processor	Allows additional CPUs to be used simultaneously by the InterBase server

A note on simultaneous connections

The USERS component of an InterBase license string indicates the maximum number of clients that can simultaneously connect to databases on the server host. The number following the word USERS determines the number of distinct users that can connect at one time. Each user can have up to four connections.

The lines in the license file must be unique. InterBase ignores duplicate instances of a certificate ID.

Tip In Borland Delphi and other VCL-based Borland tools, each *TTable* and *TQuery* component establishes a separate connection to the database by default, and therefore counts toward the connection limit. You should instead use a *TDatabase* component to connect to the database and associate each *TTable* and *TQuery* with the *TDatabase*. This way, the application uses only one database connection per *TDatabase*.

The InterBase license file

Information about what functionality has been enabled is stored in the *ib_license.dat* file in the InterBase install directory on both Windows and UNIX. IBConsole's Add Certificate function (Windows) and the *iblicense* utility (all platforms) both write to this file. Users do not have to write to this file directly. This section describes the contents of this file.

Elements in the ib_license.dat file

Each line of *ib_license.dat* contains four or more components, listed in the table below:

Table 7.4 InterBase components

Component	Description
COMMENT *string*	Descriptive comments (optional)
ID *string*	Unique identification for a key
KEY *string*	Encoded key you need to activate a server capability
OPTIONS *string*	Set of single-character codes for enabling specific InterBase server or client functionality; see below
PRODUCT *name*	INTERBASE
UNTIL *date*	Expiration date of the license (optional)
USERS *n*	User limit for the attachment governor; see below
VERSION *string*	Code for InterBase platform; for example, "WI" for Windows, "SO" for Solaris, and "LI" for Linux

The KEY component is a form of password, encoded from the combination of all the other components present in the license key string. You need the KEY in order to activate the capabilities specified in OPTIONS. Any change to the components, such as adding to the OPTIONS or altering the number of USERS, invalidates the KEY because the KEY is based on the specific values of the components.

Options in the *ib_license.dat* file

The capabilities of an InterBase installation are determined by the options present in the license file. Each license line in the file contains some number of single-letter option codes in the OPTIONS component of the line. The options and the respective functions they enable are listed below:

Table 7.5 InterBase license options

Option	Enables functionality
D	Metadata changes with CREATE, ALTER or DROP statements
E	External table access
I	Internal table access
J	For Type 3 InterClient only: Listening server capability (like the "S" option) for the InterServer component of InterClient; appears in InterServer's license file, not InterBase's
Q	Query tool (**isql**)
R	Client capability; required for clients
S	Server capability; required for servers
W	Access license with unlimited users

Chapter

8

Database Backup and Restore

A database *backup* saves a database to a file on a hard disk or other storage medium. To protect a database from power failure, disk crashes, or other potential data loss, you should regularly back up the database. For additional safety, it is recommended to store the backup medium in a different physical location from the database server.

A database *restore* re-creates a database from a backup file.

Benefits of backup and restore

Operating systems usually include facilities to archive database files. Using the InterBase backup and restore feature in **gbak** or IBConsole offers several advantages over such backup methods. The backup and restore process accomplishes the following:

- Improves database performance, by performing garbage collection on outdated records, and balances indexes.
- Reclaims space occupied by deleted records, and pack the remaining data. This often reduces database size.
- Gives you the option of changing the database page size or distributing the database among multiple files or disks.
- Enables backups to run concurrently while other users are using the database. You do not have to shut down the database to run a back up. However, any data changes that clients commit to the database after the backup begins are not recorded in the backup file.
- Provides you with a platform-independent, stable snapshot of the database for archiving purposes.

- Creates a database backup to a disk file or to a named tape device.
- Upgrades the ODS

New major releases of the InterBase server often contain changes to the on-disk structure (ODS). If the ODS has changed and you want to take advantage of any new InterBase features, upgrade your databases to the new ODS.

You need not upgrade databases to use a new version of InterBase. The new versions can still access databases created with a previous version, but cannot take advantage of any new InterBase features.

To upgrade existing databases to a new ODS, perform the following steps:

1. Before installing the new version of InterBase, back up databases using the old version.
2. Install the new InterBase server.
3. Once the new server is installed, restore the databases with the new **gbak**.

Database ownership

Although backing up a database can be performed by only the owner or SYSDBA, any user can restore a database as long as they are not restoring it over an existing database. A restored database file belongs to the user ID of the person who performed the restore. This means that backing up and restoring a database is a mechanism for changing the ownership of a database. It also means that an unauthorized user can "steal" a database by restoring a backup file to a machine where he knows the SYSDBA password. It is important to ensure that your backup files are secured from unauthorized access.

Note To restore a database over an existing database, you must be the owner of the existing database or SYSDBA.

Backing up a database using IBConsole

Use the Database Backup dialog to back up a database. To access this dialog, select a logged-in server from the list of available servers displayed in the Tree pane and continue with any of the following methods:

- Select Databases or any database under the Databases hierarchy and choose **Database | Maintenance | Backup/Restore | Backup**.
- Right-click any connected database under the Databases hierarchy and choose **Backup/Restore | Backup** from the context menu.
- Select a connected database under the Databases hierarchy and double-click Database Backup in the Work pane.

- Select a database alias under Backup in the Tree pane and double-click Backup in the Work pane.

The Database Backup dialog appears:

Figure 8.1 Database backup dialog

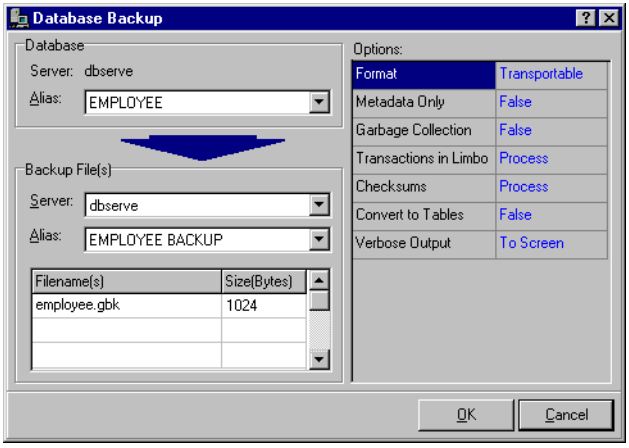

To back up a database:

1 Check the database server to make sure the server indicated is correct. If it is not, cancel this dialog and re-initiate the Database Backup dialog under the correct server.

2 If you accessed the Database Backup dialog from a database alias then the Alias field is automatically assigned. If, however, you accessed the Database Backup dialog from Databases, then you must select an alias from the list of database aliases.

Note The database alias references the necessary database files information associated with the database so you only need to specify the alias name, not the actual database filename, when indicating the database to backup. If the database spans multiple files, the server uses the header page of each file to locate additional files, so the entire database can be backed up based on the alias filename.

3 Select a destination server from a list of registered servers in the Backup Files Server drop down list.

4 Once a destination server has been selected a list of backup file aliases is available from the Backup Files Alias drop down list. If you want to overwrite an existing backup file, select the appropriate file from the drop down list. If you want to create a new backup file you can type a new alias name in the Backup File(s) Alias field.

5 Next you must indicate where the backup is to be stored by entering one or more filenames, specifying a size for each file, in the Backup File(s) table. To insert a new row into the Backup File(s) table, move to the last row and column in the table and type Ctrl-Tab.

When entering a filename, make sure to include the file path unless you wish to default the file to the working directory.

If you select an existing backup alias the table displays all the filenames and file sizes of that alias. You can edit any information within this table. You can add another file to the backup file list by entering a new filename at the end of the table. You can remove a file from the backup file list by deleting (blanking out) the values in the table.

6 You can specify backup options by entering a valid value, by clicking the option value and choosing a new value from a drop down list of values or by double-clicking the option value to rotate its value to the next in the list of values. See "Backup options" below for descriptions of these options.

7 Click OK to start the backup.

Note Database files and backup files can have any name that is legal on the operating system; the *gdb* and *gbk* file extensions are InterBase conventions only. Because files that have the *gdb* extension automatically get backed up whenever they are touched in some versions of Windows XP and ME, InterBase now recommends using an *ib* extension for database files and *ibk* for backup files.

A backup file typically occupies less space than the database because it includes only the current version of data and incurs less overhead for data storage. A backup file also does not contain index data structures, only the index definition.

If you specify a backup file that already exists, IBConsole overwrites it. To avoid overwriting, specify a unique name for the backup file.

Backup options

The backup options are shown on the right side of the Database Backup dialog. You can specify options by entering a value, by clicking the option value and choosing a new value from a drop down list of values or by double-clicking the option value to rotate its value to the next in the list of values.

Figure 8.2 Database backup options

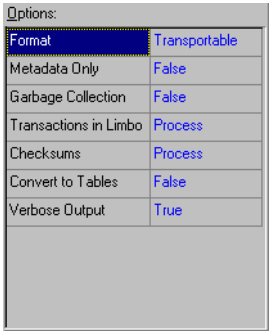

Format

Option values are Transportable and Non-transportable.

To move a database to a machine with a different operating system from the machine on which the backup was performed, make sure the Format option is set to Transportable. This option writes data in a generic format, enabling you to move to any machine that supports InterBase.

Note Never copy a database from one location to another. Back it up and then restore it to the new location.

Metadata Only

Option values are True and False.

When backing up a database, you can exclude its data, saving only its metadata. You might want to do this to:

- Retain a record of the metadata before it is modified.
- Create an empty copy of the database. The copy has the same metadata but can be populated with different data.

To back up metadata only, select True for the Metadata Only option.

Tip You can also extract a database's metadata using **isql**. **isql** produces a SQL data definition (text) file containing SQL commands. IBConsole backup Metadata Only creates a binary backup file containing only metadata.

This function corresponds to the **-metadata** option of **gbak.**

Garbage collection

Option values are True and False.

By default, IBConsole performs garbage collection during backup. To prevent garbage collection during a backup, set the Garbage Collection option value to False.

Garbage collection marks space used by old versions of data records as free for reuse. Generally, you want IBConsole to perform garbage collection during backup.

Tip You do not want to perform garbage collection if there is data corruption in old record versions and you want to prevent InterBase from visiting those records during a backup.

This function corresponds to the **-garbage_collect** option of **gbak**.

Transactions in limbo

Option values are Process and Ignore.

To ignore limbo transactions during backup, set the Transactions in Limbo option value to Ignore.

When IBConsole ignores limbo transactions during backup, it ignores all record versions created by any limbo transaction, finds the most recently committed version of a record, and backs up that version.

Limbo transactions are usually caused by the failure of a two-phase commit. They can also exist due to system failure or when a single-database transaction is prepared.

Before backing up a database that contains limbo transactions, it is a good idea to perform transaction recovery, by choosing **Database | Maintenance | Transaction Recovery** in the Database Maintenance window. Refer to "Recovering transactions" on page 6-34 for more information.

This function corresponds to the **-limbo** option of **gbak**.

Checksums

Option values are Process and Ignore.

To ignore checksums during backup, set the Checksums option value to Ignore.

A checksum is a page-by-page analysis of data to verify its integrity. A bad checksum means that a data page has been randomly overwritten; for example, due to a system crash.

Checksum errors indicate data corruption, and InterBase normally prevents you from backing up a database if bad checksums are detected. Examine the data the next time you restore the database.

This function corresponds to the **-ignore** option of **gbak**.

Convert to Tables

To convert external files to internal tables, set the Convert to Tables option value to True.

This function corresponds to the **-convert** option of **gbak**.

Verbose Output

Option values are None, To Screen and To File.

To monitor the backup process as it runs, set the Verbose Output option value to To Screen. This option opens a standard text display window to display status messages during the backup. For example:

Figure 8.3 Database backup verbose output

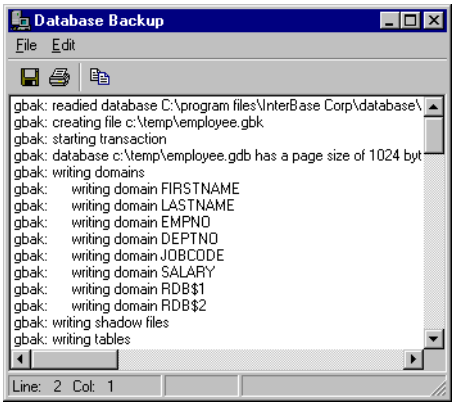

The standard text display window enables you to search for specific text, save the text to a file, and print the text. For an explanation of how to use the standard text display window, see "Standard text display window" on page 2-7.

This function corresponds to the **-verbose** option of **gbak.**

Transferring databases to servers running other operating systems

1. Set the Format option to Transportable in the Database Backup dialog.
2. Back up the database.
3. If you backed up to a removable medium, proceed to Step 4. If you created a backup file on disk, use operating-system commands to copy the file to a removable medium, such as a tape. Then load the contents of the medium onto another machine, or copy it across a network to another machine.
4. On the destination machine, restore the backup file. If restoring from a removable medium, such as tape, specify the device name instead of the backup file.

Restoring a database using IBConsole

Use the Database Restore dialog to restore databases. To access this dialog, select a server from the list of available servers displayed in the Tree pane and continue with one of these possible methods:

- Select anything under the databases hierarchy and choose **Database | Maintenance | Backup/Restore | Restore**.
- Double-click any backup alias name under the Backup hierarchy.
- Right-click Backup or any backup alias name under the Backup hierarchy and choose Restore from the context menu.
- Select any backup alias name under Backup and click Restore in the Work pane.

The Database Restore dialog appears:

Figure 8.4 Database Restore dialog

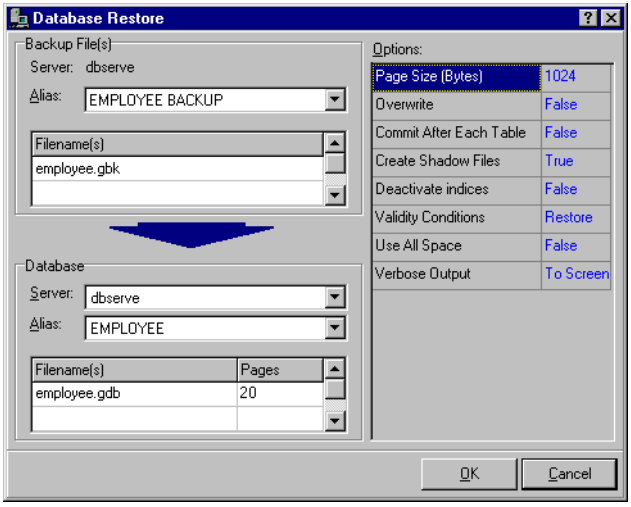

Important When restoring a database, do not replace a database that is currently in use.

To restore a database:

1. Check the source Backup File(s) Server to make sure the server indicated is correct. If it is not, cancel this dialog and re-initiate the Database Restore dialog under the correct server.
2. If you accessed the Database Restore dialog from a backup alias, then the Alias field is automatically assigned. If, however, you accessed the Database Restore dialog from Backup, then you must select an alias from the list of backup aliases.

Note The backup alias references the necessary database backup files information associated with the database so you only need to specify the alias name, not the actual backup filename, when indicating the backup to restore. If the backup spans multiple files, the server uses header page of each file to locate additional files, so the entire backup can be restored based on the alias filename.

3 If you choose a backup file alias, the Backup File(s) table displays the associated backup files. If you do not specify a backup file alias, then you can enter the backup filenames manually, or you can browse for the file by selecting "File..." from the Alias drop-down list. When entering a filename, be sure to include the file path where the file is located. It is important that you include all filenames associated with the restore. To insert a new row into the Backup File(s) table, move to the last row and column in the table and type Ctrl-Tab.

4 Select a destination server from a list of registered servers in the Database Server drop down list.

5 Once a destination server has been selected a list of database aliases is available from the Database Alias drop down list. If you want to restore to an existing database, select the appropriate alias from the drop down list. If you want to restore to a new database, type a new alias name in the Database Alias field.

6 Next you must indicate where the backup is to be restored to by entering one or more filenames, specifying the number of pages required for each file, in the Database table. When entering a filename, make sure to include the file path unless you wish to default the file to the working directory. To insert a new row into the Database table, move to the last row and column in the table and type Ctrl-Tab.

   You might want to restore a database to multiple files to distribute it among different disks, which provides more flexibility in allocating system resources.

   You can add another file to the backup file list by entering a new filename at the end of the table. You can remove a file from the backup file list by deleting (blanking out) the values in the table.

   If you selected an existing database alias the Database table will display all the associated filenames and number of pages. You can edit any information within this table. You can add another file to the database file list by entering a new filename at the end of the table. You can remove a file from the list by deleting (blanking out) the values in the table.

   Note You cannot restore a database to a network file system (mapped drive).

7 You can specify restore options by entering a valid value, by clicking the option value and choosing a new value from a drop down list of values or by double-clicking the option value to rotate its value to the next in the list of values. See "Restore options" below for a description of these options.

8 Click OK to start the restore.

Typically, a restored database occupies less disk space than it did before being backed up, but disk space requirements could change if the ODS version changes. For information about the ODS, see "Benefits of backup and restore" on page 8-1.

Note InterBase 6 restore allows you to restore a database successfully even if for some reason the restore process could not rebuild indexes for the database. For example, this can occur if there is not enough temporary disk space to perform the sorting task necessary to build an index during the restore. If this occurs, the database is restored and available, but indexes are inactive. This is as if you had set the Deactivate Indexes option to True, or used the **-i** switch of **gbak**. After the restore completes, use ALTER INDEX to set the indexes active.

Restore options

The restore options are shown on the right side of the Database Restore dialog. You can specify options by entering a value, by clicking the option value and choosing a new value from a drop down list of values or by double-clicking the option value to rotate its value to the next in the list of values.

Figure 8.5 Database restore options

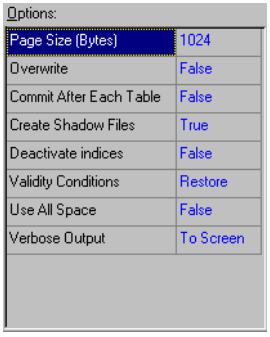

Page Size

InterBase supports database page sizes of 1024, 2048, 4096, and 8192 bytes. The default is 4096 bytes. To change page size, back up the database and then restore it, modifying the Page Size option in the Database Restore dialog.

Changing the page size can improve performance for the following reasons:

- Storing and retrieving Blob data is most efficient when the entire Blob fits on a single database page. If an application stores many Blobs exceeding 4KB, using a larger page size reduces the time for accessing Blob data.
- InterBase performs better if rows do not span pages. If a database contains long rows of data, consider increasing the page size.

- If a database has a large index, increasing the database page size reduces the number of levels in the index tree. Indexes work faster if their depth is kept to a minimum. Choose **Database | Maintenance | Database Statistics** to display index statistics, and consider increasing the page size if index depth is greater than three on any frequently used index.
- If most transactions involve only a few rows of data, a smaller page size may be appropriate, because less data needs to be passed back and forth and less memory is used by the disk cache.

This function corresponds to the **-page_size** option of **gbak**.

Overwrite

Option values are True and False.

IBConsole cannot overwrite an existing database file unless the Overwrite option value is set to True. If you attempt to restore to an existing database name and this option is set to False, the restore does not proceed.

Note To restore a database over an existing database, you must be the owner of the existing database or SYSDBA.

Important Do not replace an existing database while clients are operating on it. When restoring to an existing file name, a safer approach is to rename the existing database file, restore the database, then drop or archive the old database as needed.

This function corresponds to the **-replace** option of **gbak**.

Commit After Each Table

Option values are True and False.

Normally, IBConsole restores all metadata before restoring any data. If you set the Commit After Each Table option value to True, IBConsole restores the metadata and data for each table together, committing one table at a time.

This option is useful when you are having trouble restoring a backup file; for example, if the data is corrupt or invalid according to integrity constraints.

If you have a problem backup file, restoring the database one table at a time lets you recover some of the data intact. You can restore only the tables that precede the bad data; restoration fails the moment it encounters bad data.

This function corresponds to the **-one_at_a_time** option of **gbak**.

Create Shadow Files

Shadow files are identical, physical copies of database files in a database. To recreate shadow files that were saved during the backup process set the Create Shadow Files option to True. For further information on shadowing see "Shadowing" on page 6-12.

Deactivate Indexes

Option values are True and False.

Normally, InterBase rebuilds indexes when a database is restored. If the database contained duplicate values in a unique index when it was backed up, restoration fails. Duplicate values can be introduced into a database if indexes were temporarily made inactive (for example, to allow insertion of many records or to rebalance an index).

To enable restoration to succeed in this case, set the Deactivate Indexes option to True. This makes indexes inactive and prevents them from rebuilding. Then eliminate the duplicate index values, and re-activate indexes through ALTER INDEX in **isql**.

A unique index cannot be activated using the ALTER INDEX statement; a unique index must be dropped and then created again. For more information about activating indexes, see the *Language Reference*.

Tip The Deactivate Indexes option is also useful for bringing a database online more quickly. Data access is slower until indexes are rebuilt, but the database is available. After the database is restored, users can access it while indexes are reactivated.

This function corresponds to the **-inactive** option of **gbak**.

Validity Conditions

Option values are Restore and Ignore.

If you redefine validity constraints in a database where data is already entered, your data might no longer satisfy the validity constraints. You might not discover this until you try to restore the database, at which time an error message about invalid data appears.

Important Always make a copy of metadata before redefining it; for example, by extracting it using **isql**.

To restore a database that contains invalid data, set the Validity Conditions option to Ignore. This option deletes validity constraints from the metadata. After the database is restored, change the data to make it valid according to the new integrity constraints. Then add back the constraints that were deleted.

This option is also useful if you plan to redefine the validity conditions after restoring the database. If you do so, thoroughly test the data after redefining any validity constraints.

This function corresponds to the **-no_validity** option of **gbak**.

Use All Space

Option values are True and False.

To restore a database with 100% fill ratio on every data page, instead of the default 80% fill ratio, set the Use All Space option to True.

This function corresponds to the **-use_all_space** option of **gbak**.

Verbose Output

Option values are None, To Screen, and To File.

To monitor the restore process as it runs, set the Verbose Output option to To Screen. This option opens a standard text display window to display status messages during the restore. For example:

Figure 8.6 Database restore verbose output

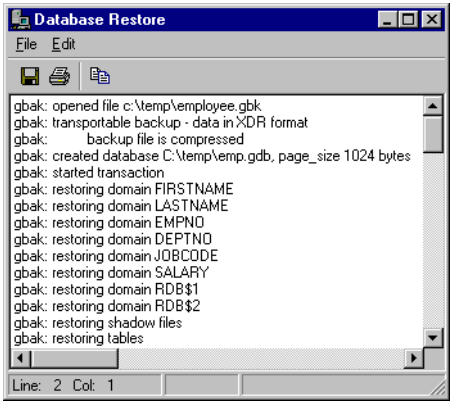

The standard text display window enables you to search for specific text, save the text to a file, and print the text. For an explanation of how to use the standard text display window, see "Standard text display window" on page 2-7.

This function corresponds to the **-verbose** option of **gbak**.

gbak command-line tool

The **gbak** command-line tool allows both back up or restore of a database, with options for changing specified database characteristics. Only SYSDBA or database owner can back up a database.

Database backup

Syntax **For backing up to a single file:**

```
gbak [-B] [options] database target
```

When backing up from a multifile database, specify only the first file name.

For backing up to multiple files:

```
gbak [-B] [options] database target1 size1[k|m|g] target2
   [size2[k|m|g] target3
```

When backing up from a multifile database, specify only the first file name.

Table 8.1 gbak arguments

Argument	Description
database	• Name of a database to back up • For a multifile database, the name of the first database file
target	Name of a storage device or backup file to which to back up • On UNIX, can also be *stdout*, in which case **gbak** writes its output to the standard output (usually a pipe) • No size need be specified when restoring to a single file, since the database always expands as needed to fill all available space
size	Length of a backup file or restored database file • The only permissible unit for a restored database file is database pages; minimum value is 200 • Default unit for a backup file is bytes • Size of backup files can also be specified in kilobytes, megabytes, or gigabytes • Do not specify a size for the final backup file or database file; the last file always expands as needed to fill all available space

Options In the OPTION column of the following tables, only the characters outside the square brackets ([ ]) are required.

Table 8.2 lists the options to **gbak** that are available for creating backups.

Table 8.2 gbak backup options

Option	Description
-b[ackup_database]	Backs up database to file or device
-co[nvert]	Converts external files as internal tables
-e[xpand]	Creates a noncompressed back up
-fa[ctor] ***n***	Uses blocking factor *n* for tape device
-g[arbage_collect]	Does not garbage collect during backup
-ig[nore]	Ignores checksums during backup
-l[imbo]	Ignores limbo transactions during backup
-m[etadata]	Backs up metadata only, no data
-nt	Creates the backup in nontransportable format
-ol[d_descriptions]	Backs up metadata in old-style format

Table 8.2 gbak backup options (*continued*)

Option	Description
-pas[sword] *text*	Checks for password *text* before accessing a database
-role *name*	Connects as role *name*
-se[rvice] *servicename*	• Creates the backup files on the host where the original database files are located, using InterBase's Service Manager • *servicename* invokes the Service Manager on the server host; syntax varies with the network protocol in use: TCP/IP `hostname:service_mgr` SPX `hostname@service_mgr` Named pipes `\\hostname\service_mgr` Local `service_mgr`
-t[ransportable]	Creates a transportable backup [default]
-u[ser] *name*	Checks for user *name* before accessing remote database
-v[erbose]	Shows what **gbak** is doing
-y [*file* | **suppress_output]**	Direct status messages to *file*; *file* must not already exist; **suppress_output** suppress output messages
-z	Show version of **gbak** and of InterBase engine

Backing up a database with gbak

When backing up a database, bear the following points in mind:

- Only the database owner or SYSDBA can back up a database.
- Unless the **-service** option is specified, **gbak** writes the backup files to the current directory of the machine on which it is running, not on the server where the database resides. If you specify a location for the backup file, it is relative to the machine where **gbak** is executing. You can write the backup files only to this local machine or to drives that are mapped to it. Note that the **-service** switch changes this behavior. (See "Using gbak with InterBase Service Manager" on page 8-19.)
- When you are backing up a multifile database, specify *only* the first file in the backup command. You must not name the subsequent database files: they will be interpreted as backup file names.

- The default unit for backup files is bytes. You can choose to specify kilobytes, megabytes, or gigabytes (**k**, **m**, or **g**) instead. Restored database files can be specified only in database pages.

Note It is good security practice to change your backup files to read-only at the system level after creating them. This prevents them from being accidentally overwritten. In addition, you can protect your databases from being "kidnapped" on UNIX and NT/2000 systems by placing the backup files in directories with restricted access.

Tip Use the **-transportable** switch if you operate in a multiplatform environment. This switch permits the database to be backed up to a platform other than the one on which it originally resided. Using this option routinely is a good idea when you are operating in a multiplatform environment.

Tip Use the **-service** switch if you are backing up to the same server that holds the original database. This option invokes the InterBase Service Manager on the server host and saves both time and network traffic.

Restoring a database with gbak

Syntax **For restoring:**

```
gbak {-C|-R} [options] source dbfile
```

For restoring to multiple files:

```
gbak {-C|-R} [options] source dbfile1 size1 dbfile2
   [size2 dbfile3 …]
```

For restoring from multiple files:

```
gbak {-C|-R} [options] source1 source2 [source3 …] dbfile
```

By extension, you can restore *from* multiple files *to* multiple files using the following syntax:

```
gbak {-C|-R} [options] source1 source2 [source3 …] dbfile1 size1
   dbfile2 [size2 dbfile3 …]
```

Table 8-7:

Argument	Description
source	Name of a storage device or backup file from which to restore On UNIX, this can also be *stdin*, in which case **gbak** reads input from the standard input (usually a pipe).
dbfile	The name of a restored database file
size	Length of a backup file or restored database file • The only permissible unit for a restored database file is database pages; minimum value is 200 • Default unit for a backup file is bytes • Size of backup files can also be specified in kilobytes, megabytes, or gigabytes • Do not specify a size for the final backup file or database file; the last file always expands as needed to fill all available space

Table 8.1 lists **gbak** options that are available when restoring databases.

Table 8.1 gbak restore options

Option	Description
-c[reate_database]	Restores database to a new file
-bu[ffers]	Sets cache size for restored database
-i[nactive]	Makes indexes inactive upon restore
-k[ill]	Does not create any shadows that were previously defined
-mo[de] {read_write | read_only}	Specifies whether the restored database is writable • Possible values are **read_only** and **read_write** • Default is **read_write**
-n[o_validity]	Deletes validity constraints from restored metadata; allows restoration of data that would otherwise not meet validity constraints
-o[ne_at_a_time]	Restores one table at a time; useful for partial recovery if database contains corrupt data
-p[age_size] ***n***	Resets page size to *n* bytes (1024, 2048, 4096, or 8192); default is 4096
-pa[ssword] ***text***	Checks for password *text* before accessing a database
-r[eplace_database]	Restores database to new file or replaces existing file

Table 8.1 gbak restore options (*continued*)

Option	Description
-se[rvice] ***servicename***	• Creates the restored database on the host where the backup files are located, using InterBase's Service Manager • *servicename* invokes the Service Manager on the server host; syntax varies with the network protocol in use: TCP/IP `hostname:service_mgr` NetBEUI `\\hostname\service_mgr` Local `service_mgr`
-user ***name***	Checks for user *name* before accessing database
-use_[all_space]	Restores database with 100% fill ratio on every data page, instead of the default 80% fill ratio
-v[erbose]	Shows what **gbak** is doing
-y [*file* | **suppress_output]**	If used with **-v**, directs status messages to *file;* if used without **-v** and *file* is omitted, suppresses output messages
-z	Show version of **gbak** and of InterBase engine

When restoring a database, bear the following points in mind:

- Anyone can restore a database. However, only the database owner or SYSDBA can restore a database over an existing database.
- Do not restore a backup over a database that is currently in use; it is likely to corrupt the database.
- When restoring from a multifile backup, name all the backup files, in any order.
- Do not provide a file size for the last (or only) file of the restored database. InterBase does not return an error, but it always "grows" the last file as needed until all available space is used. This dynamic sizing is a feature of InterBase.
- You specify the size of a restored database in database pages. The default size for database files is 200 pages. The default database page size is 4K, so if the page size has not been changed, the default database size is 800K. This is sufficient for only a very small database. To change the size of the database pages, use the **-p[age_size]** option when restoring.

Tip Use the **-service** switch if you are restoring to the same server that holds the backup file. This option invokes the InterBase Service Manager on the server host and saves both time and network traffic.

Note If you specify several target database files but have only a small amount of data, the target files are quite small (around 800K for the first one and 4K for subsequent ones) when they are first created. They grow in sequence to the specified sizes as you populate the database.

Using gbak with InterBase Service Manager

When you run **gbak** with the **-service** switch, **gbak** invokes the backup and restore functions of InterBase's Service Manager on the server where the database resides. When run without the **-service** switch, **gbak** executes on the machine where it is invoked—typically a client—and writes the backup file on (or relative to) that machine. Using the **-service** switch to invoke the Service Manager saves a significant amount of time and network traffic when you want to create the backup on the same host on which the database resides. You have the option of specifying another machine as the target when using the -**service** switch, but the advantages of reduced time and network traffic are lost.

When you use the **-service** switch, you specify the host name followed by the string "service_mgr". The syntax you use for this varies with the network protocol you are using. Together, these components are referred to as "host_service" in the syntax statements that follow in this section.

Table 8.2 *host_service* syntax for calling the Service Manager with **gbak**

Network protocol	Syntax
TCP/IP	*hostname*:service_mgr
NetBEUI	*hostname*\service_mgr
Local	service_mgr

The syntax in the right column appears in the **gbak** syntax below as "host_service."

The "local" case is trivial on NT. If you are backing up a local database, the results in terms of time and network traffic are the same whether you use the **-service** switch or not, even though the actual implementation would be slightly different. On UNIX systems, the local case is equivalent to specifying (for TCP/IP) **localhost:service_mgr** and saves both time and network traffic.

Syntax **Backing up with Service Manager**

```
gbak -b [options] -se[rvice] host_service database filename
```

Syntax **Restoring with Service Manager**

```
gbak {-c|-r} [options] -se[rvice] host_service filename database
```

You can back up to multiple files and restore from multiple files using Service Manager.

Important On UNIX systems, in order to restore a database that has been backed up using the Service Manager, you must either use the Service Manager for the restore or you must be logged onto the system as the user that InterBase was running as when the backup was created (either *root* or *interbase*). This is because the InterBase user (*root* or *interbase*) is the owner of the backup file at the system level when the Service Manager is invoked, and the backup file is readable to only that user. When **gbak** is used to back up a database without the **-service** option, the owner of the backup file at the system level is the login of the person who ran **gbak**. On Windows platforms, the system-level constraints do not apply.

The user name and password

When InterBase checks to see whether the user running **gbak** is authorized to do so, it determines the user according to the following hierarchy:

- The -**user** that is specified, with a correct password, as part of the **gbak** command
- The user and password specified in the ISC_USER and ISC_PASSWORD environment variables, provided they also exist in the InterBase security database. (Setting these environment variables is strongly *not* recommended, since it is extremely insecure.)
- UNIX only: If no user is specified at any of the previous levels, InterBase uses the UNIX login if the user is running on the server or on a trusted host.

Some backup and restore examples

Note The following examples use forward slashes exclusively. InterBase accepts either forward or backward slashes for paths on Wintel platforms.

Backup examples

The following example backs up *foo.ib*, which resides on the server *jupiter* and writes the backup file to the current directory of the client machine where **gbak** is running. *foo.ib* can be either a single-file database or the name of the first file in a multifile database. Using this syntax (without the **-se** switch) copies a lot of data over the net.

```
gbak -b -user joe -pa blurf@ jupiter:/foo.ib foo.ibk
```

The next example backs up *foo.ib*, which resides on the server *jupiter* and writes the backup file to the *C:/archive* directory on the client machine where **gbak** is running. As before, *foo.ib* can be a single file database or the name of the first file in a multifile database. This syntax causes the same amount of network traffic as the first example.

```
gbak -b -user joe -pa blurf@ jupiter:/foo.ib C:\archive\foo.ibk
```

The next example backs up the same database on *jupiter*, but uses the **-se**[**rvice**] switch to invoke the Service Manager on *jupiter*, which writes the backup to the *backup* directory on *jupiter*. This command causes very little network traffic and is therefore faster than performing the same task without the **-se** (**-service**) switch. Note that the syntax (`jupiter:service_mgr`) indicates a TCP/IP connection.

```
gbak -b -user joe -pa blurf@ -se jupiter:service_mgr
   /foo.ib /backup/foo.ibk
```

The next example again backs up *foo1.ib* on server *jupiter* to multiple files in the */backup* directory on *jupiter* using the Service Manager. This syntax backs up a single file or multifile database and uses a minimum of time and network traffic. It converts external files as internal tables and creates a backup in a transportable format that can be restored on any InterBase-supported platform. To back up a multifile database, name only the first file in the backup command. In this example, the first two backup files are limited to 500K. The last one expands as necessary.

```
gbak -b -user joe -pa blurf@ -co -t -se jupiter:service_mgr
   /foo1.ib /backup/backup1.ibk 500k /backup/backup2.ibk 500k
   /backup/lastBackup.ibk
```

Database restore examples

The first example restores a database that resides in the */archive* directory on the machine where **gbak** is running and restores it to *jupiter*, overwriting an existing (but inactive) database.

```
gbak -r -user joe -pa blurf@ C:\archive\foo.ibk jupiter:/foo.ib
```

The next example restores a multifile database from the */backup* directory of jupiter to the */companydb* directory of *jupiter*. This command runs on the server by invoking Service Manager, thus saving time and network traffic. In this example, the first two files of the restored database are 500 pages long and the last file grows as needed.

```
gbak -r user -joe -pa blurf@ -se jupiter:service_mgr /backup/foo1.ibk
   /backup/foo2.ibk /backup/fooLast.ibk /companydb/foo1.ib 500
   /companydb/foo2.ib 500 /companydb/fooLast.ib
```

The next example executes on server Jupiter using Service Manager and restores a backup that is on Jupiter to another server called Pluto.

```
gbak -r user -joe -pa blurf@ -se jupiter:service_mgr
   /backup/foo.ibk pluto:/companydb/foo.ib
```

gbak error messages

Table 8.3 gbak backup and restore error messages

Error Message	Causes and Suggested Actions to Take
Array dimension for column <string> is invalid	Fix the array definition before backing up
Bad attribute for RDB$CHARACTER_SETS	An incompatible character set is in use
Bad attribute for RDB$COLLATIONS	Fix the attribute in the named system table
Bad attribute for table constraint	Check integrity constraints; if restoring, consider using the **-no_validity** option to delete validity constraints
Blocking factor parameter missing	Supply a numeric argument for "factor" option
Cannot commit files	• Database contains corruption or metadata violates integrity constraints • Try restoring tables using **-one_at_a_time** option, or delete validity constraints using **-no_validity** option
Cannot commit index <string>	• Data might conflict with defined indexes • Try restoring using "inactive" option to prevent rebuilding indexes
Cannot find column for Blob	
Cannot find table <string>	
Cannot open backup file <string>	Correct the file name you supplied and try again
Cannot open status and error output file <string>	• Messages are being redirected to invalid file name • Check format of file or access permissions on the directory of output file
Commit failed on table <string>	• Data corruption or violation of integrity constraint in the specified table • Check metadata or restore "one table at a time"
Conflicting switches for backup/restore	A backup-only option and restore-only option were used in the same operation; fix the command and execute again
Could not open file name <string>	Fix the file name and re-execute command
Could not read from file <string>	Fix the file name and re-execute command
Could not write to file <string>	Fix the file name and re-execute command
Datatype n not understood	An illegal datatype is being specified

Table 8.3 gbak backup and restore error messages (*continued*)

Error Message	Causes and Suggested Actions to Take
Database format n is too old to restore to	• The **gbak** version used is incompatible with the InterBase version of the database • Try backing up the database using the **-expand** or **-old** options and then restoring it
Database <string> already exists. To replace it, use the -R switch	• You used **-create** in restoring a back up file, but the target database already exists • Either rename the target database or use **-replace**
Could not drop database <string> (database might be in use).	• You used **-replace** in restoring a file to an existing database, but the database is in use • Either rename the target database or wait until it is not in use
Device type not specified	The **-device** option (Apollo only) must be followed by **ct** or **mt**; obsolete as of InterBase V3.3
Device type <string> not known	The **-device** option (Apollo only) was used incorrectly; obsolete as of InterBase V3.3
Do not recognize record type n	
Do not recognize <string> attribute n -- continuing	
Do not understand BLOB INFO item n	
Error accessing BLOB column <string> -- continuing	
ERROR: Backup incomplete	• The backup cannot be written to the target device or file system • Either there is insufficient space, a hardware write problem, or data corruption
Error committing metadata for table <string>	• A table within the database could be corrupt. • If restoring a database, try using **-one_at_a_time** to isolate the table
Exiting before completion due to errors	• This message accompanies other error messages and indicates that back up or restore could not execute • Check other error messages for the cause.
Expected array dimension n but instead found m	Try redefining the problem array
Expected array version number n but instead found m	Try redefining the problem array

Table 8.3 **gbak** backup and restore error messages (*continued*)

Error Message	Causes and Suggested Actions to Take
Expected backup database <string>, found <string>	Check the name of the backup file being restored
Expected backup description record	
Expected backup start time <string>, found <string>	
Expected backup version 1, 2, or 3. Found n	
Expected blocking factor, encountered <string>	The **-factor** option requires a numeric argument
Expected data attribute	
Expected database description record	
Expected number of bytes to be skipped, encountered <string>	
Expected page size, encountered <string>	The **-page_size** option requires a numeric argument
Expected record length	
Expected volume number n, found volume n	When backing up or restoring with multiple tapes, be sure to specify the correct volume number
Expected XDR record length	
Failed in put_blr_gen_id	
Failed in store_blr_gen_id	
Failed to create database <string>	The target database specified is invalid; it might already exist
column <string> used in index <string> seems to have vanished	• An index references a non-existent column • Check either the index definition or column definition
Found unknown switch	An unrecognized **gbak** option was specified
Index <string> omitted because n of the expected m keys were found	
Input and output have the same name. Disallowed.	A backup file and database must have unique names; correct the names and try again
Length given for initial file (n) is less than minimum (m)	• In restoring a database into multiple files, the primary file was not allocated sufficient space • InterBase automatically increases the page length to the minimum value • No action necessary

Table 8.3 **gbak** backup and restore error messages (*continued*)

Error Message	Causes and Suggested Actions to Take
Missing parameter for the number of bytes to be skipped	
Multiple sources or destinations specified	Only one device name can be specified as a source or target
No table name for data	• The database contains data that is not assigned to any table • Use **gfix** to validate or mend the database
Page size is allowed only on restore or create	The **-page_size** option was used during a back up instead of a restore
Page size parameter missing	The **-page_size** option requires a numeric argument
Page size specified (n bytes) rounded up to m bytes	Invalid page sizes are rounded up to 1024, 2048, 4096, or 8192, whichever is closest
Page size specified (n) greater than limit (8192 bytes)	Specify a page size of 1024, 2048, 4096, or 8192
Password parameter missing	• The back up or restore is accessing a remote machine • Use **-password** and specify a password
Protection is not there yet	Unimplemented option **-unprotected** used
Redirect location for output is not specified	You specified an option reserved for future use by InterBase
REPLACE specified, but the first file <string> is a database	Check that the file name following the **-replace** option is a backup file rather than a database
Requires both input and output file names	Specify both a source and target when backing up or restoring
RESTORE: decompression length error	• Possible incompatibility in the **gbak** version used for backing up and the **gbak** version used for restoring • Check whether **-expand** should be specified during back up
Restore failed for record in table <string>	Possible data corruption in the named table
Skipped n bytes after reading a bad attribute n	
Skipped n bytes looking for next valid attribute, encountered attribute m	
Trigger <string> is invalid	

Table 8.3 **gbak** backup and restore error messages (*continued*)

Error Message	Causes and Suggested Actions to Take
Unexpected end of file on backup file	• Restoration of the backup file failed; the backup procedure that created the backup file might have terminated abnormally • If possible, create a new backup file and use it to restore the database
Unexpected I/O error while <string> backup file	A disk error or other hardware error might have occurred during a backup or restore
Unknown switch <string>	An unrecognized **gbak** option was specified
User name parameter missing	• The backup or restore is accessing a remote machine • Supply a user name with the **-user** option
Validation error on column in table <string>	• The database cannot be restored because it contains data that violates integrity constraints • Try deleting constraints from the metadata by specifying **-no_validity** during restore
Warning -- record could not be restored	Possible corruption of the named data
Wrong length record, expected n encountered n	

Chapter

9

Database Statistics and Connection Monitoring

InterBase provides a number of ways to view statistics about database behavior and to exert control over that behavior. This chapter provides a description of the following InterBase facilities:

- Monitoring with system temporary tables
- Viewing statistics using IBConsole
- The gstat command-line tool
- Viewing lock statistics
- Retrieving statistics with api_database_info()

Monitoring with system temporary tables

The InterBase Server has always kept a lot of statistics about what was going on, but it has not been easy, or in some cases possible, to surface that information. Now, InterBase captures that information and makes it available in a set of global system temporary tables. These tables describe the runtime behavior of a database. They also provide a level of control.

Although it has always been possible to see a list of users who were currently attached to a database, you can now find out much more. For example, you can see how long each user has been connected, what application each user is running, or the total amount of data I/O used by each attachment. A glance at the temporary table metadata listed on pages 7-31 to 7-42 of the *Language Reference* will suggest the vast possibilities that are available here.

It is also possible to exercise a certain amount of control over the state of a database by performing updates to these tables. See "Updating system temporary tables" on page 9-4.

These system temporary tables are specific to each database attachment and are visible only to the sysdba user and the database owner. There is therefore no need for unique names and no danger of collisions by separate attachments. Each table is populated only at the point when a client queries it.

The following system temporary tables are available. Their structure is documented in the *Language Reference* on pages 7-31 to 7-42.

Table 9.1 InterBase temporary system tables

Table name	Description
TMP$ATTACHMENTS	One row for each connection to a database
TMP$DATABASE	One row for each database you are attached to
TMP$POOL_BLOCKS	One row for each block of memory in each pool
TMP$POOLS	One row for each current memory pool
TMP$PROCEDURES	One row for each procedure executed since the current connection began
TMP$RELATIONS	One row for each relation referenced since the current connection began
TMP$STATEMENTS	One row for each statement currently executing for any current connection
TMP$TRANSACTIONS	One row for each transaction that is active or in limbo

Querying system temporary tables

Clients can query these tables using SELECT statements, just as they would query any other table. By querying these tables, a rich collection of data about server performance and user behavior is available.

You cannot create or redefine temporary tables yourself.

Tip For frequent monitoring, the best transaction control is to start the transaction as READ_COMMITTED, READ_ONLY. Then commit it with COMMIT_RETAINING. This has the least impact on the system.

Refreshing the temporary tables

To refresh the rows in the temporary tables, commit your transaction and perform the SELECT from the temporary tables again. InterBase automatically deletes the rows stored in temporary tables on a commit.

Listing the temporary tables

To display a list of these temporary tables, issue the following command in **isql**:

```
SHOW SYSTEM
```

The temporary tables are listed at the end of the system tables. To see the metadata for a particular table, issue:

```
SHOW TABLE tablename
```

Note The SHOW SYSTEM command is available only in command-line **isql**, not in InterBase Windows **isql**.

Security

Unlike system tables, which have a default access privilege of SELECT for PUBLIC users, the temporary tables have no default access by PUBLIC. The display and manipulation of this runtime information is restricted to SYSDBA and the database owner. These two users have the option of using the GRANT statement to allow access to other users. The statement can grant only SELECT privileges.

Examples

To illustrate the richness of the possibilities afforded by these temporary tables, here are some examples how you might query them.

Top ten SQL statements by execution

```
SELECT a.tmp$user, s.tmp$timestamp, s.tmp$sql, s.tmp$quantum
   FROM TMP$STATEMENTS s, TMP$ATTACHMENTS a
   WHERE a.TMP$ATTACHMENT_ID = s.TMP$ATTACHMENT_ID
   ORDER BY s.TMP$QUANTUM DESC ROWS 10;
```

Top ten oldest transaction snapshots

```
SELECT a.TMP$USER, t.TMP$TIMESTAMP, t.TMP$TRANSACTION_ID, t.TMP$SNAPSHOT
   FROM TMP$ATTACHMENTS a, TMP$TRANSACTIONS t
   WHERE a.TMP$ATTACHMENT_ID = t.TMP$ATTACHMENT_ID
   ORDER BY t.TMP$SNAPSHOT ROWS 10;
```

Top ten tables with the most garbage to clean up

```
SELECT TMP$RELATION_NAME, TMP$GARBAGE_COLLECT_PAGES
   FROM TMP$RELATIONS
   ORDER BY TMP$GARBAGE_COLLECT_PAGES DESC ROWS 10;
```

Top ten most executed stored procedures

```
SELECT TMP$PROCEDURE_NAME, TMP$INVOCATIONS
   FROM TMP$PROCEDURES
   ORDER BY TMP$INVOCATIONS DESC ROWS 10;
```

Is database sweep active and what's its progress?

```
SELECT TMP$SWEEP_RELATION, TMP$SWEEP_RECORDS
   FROM TMP$DATABASE
   WHERE TMP$SWEEP_ACTIVE = 'Y';
```

Pool memory allocations grouped by pool type

```
SELECT TMP$TYPE, SUM(TMP$POOL_MEMORY) TMP$TOTAL_MEMORY,
   SUM(TMP$FREE_MEMORY) TMP$TOTAL_FREE
   FROM TMP$POOLS
   GROUP BY TMP$TYPE
   ORDER BY 2 DESC;
```

Updating system temporary tables

There are cases where, having acquired information about the state of the database, you need to take appropriate action. You might, for example, detect a transaction that had unexpectedly been open for many hours, or one that was consuming resources that were needed by others. By updating the TMP$STATE column of certain temporary tables, you can perform the following updates:

- Roll back an active or limbo transaction
- Commit a limbo transaction
- Cancel an attachment's executing operation
- Shut down the current attachment
- Make an executing statement stop running

To roll back an active transaction

```
UPDATE TMP$TRANSACTIONS SET TMP$STATE = 'ROLLBACK' WHERE TMP$TRANSACTION_ID=123;
```

To roll back a limbo transaction

```
UPDATE TMP$TRANSACTIONS SET TMP$STATE = 'ROLLBACK' WHERE TMP$TRANSACTION_ID=123;
```

To commit a limbo transaction

```
UPDATE TMP$TRANSACTIONS SET TMP$STATE = 'COMMIT' WHERE TMP$TRANSACTION_ID=123;
```

To cancel the attachment's currently executing operation

```
UPDATE TMP$ATTACHMENTS SET TMP$STATE = 'CANCEL' WHERE TMP$ATTACHMENT_ID=123;
```

To shut down the current attachment

```
UPDATE TMP$ATTACHMENTS SET TMP$STATE = 'SHUTDOWN' WHERE TMP$ATTACHMENT_ID=123;
```

Shutting down an attachment detaches the user from the database and terminates the local or network attachment to the server.

To make an executing statement stop running

```
UPDATE TMP$STATEMENTS SET TMP$STATE = 'CANCEL' WHERE TMP$STATEMENT_ID=123;
```

Making global changes

The above examples operate on a single attachment or transaction. You can make more global changes. For example:

To roll back all active transactions

```
UPDATE TMP$TRANSACTIONS SET TMP$STATE = 'ROLLBACK' WHERE TMP$STATE ='ACTIVE';
```

To roll back all limbo transactions

```
UPDATE TMP$TRANSACTIONS SET TMP$STATE = 'ROLLBACK' WHERE TMP$STATE ='LIMBO';
```

To commit all limbo transactions

```
UPDATE TMP$TRANSACTIONS SET TMP$STATE = 'COMMIT' WHERE TMP$STATE ='LIMBO';
```

Viewing statistics using IBConsole

To view database statistics, use one of the following methods to access the Database Statistics dialog:

- Select a connected database in the Tree pane and choose **Database | Maintenance | Database Statistics**.
- Select a connected database in the Tree pane and double-click Database Statistics in the Work pane.
- Right-click a connected database in the Tree pane and choose **Maintenance | Database Statistics** from the context menu.

A Database Statistics dialog appears where you can select which statistics you want to display.

Figure 9.1 Database Statistics options

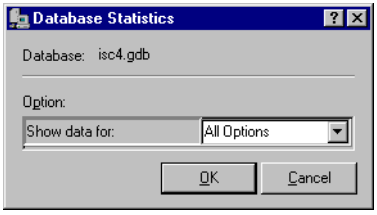

To view database statistics:

1. Select the statistical data you wish to generate from the Options list.

   You can specify options by entering a value, by clicking the option value and choosing a new value from a drop down list of values or by double-clicking the option value to rotate its value to the next in the list of values.

2. Click OK to generate database statistics.

Note In some cases, it can take a long time to display the statistics for large databases because, depending on what information has been selected to report, generating these statistics may analyze all the tables and indexes in a database.

Figure 9.2 Database Statistics dialog

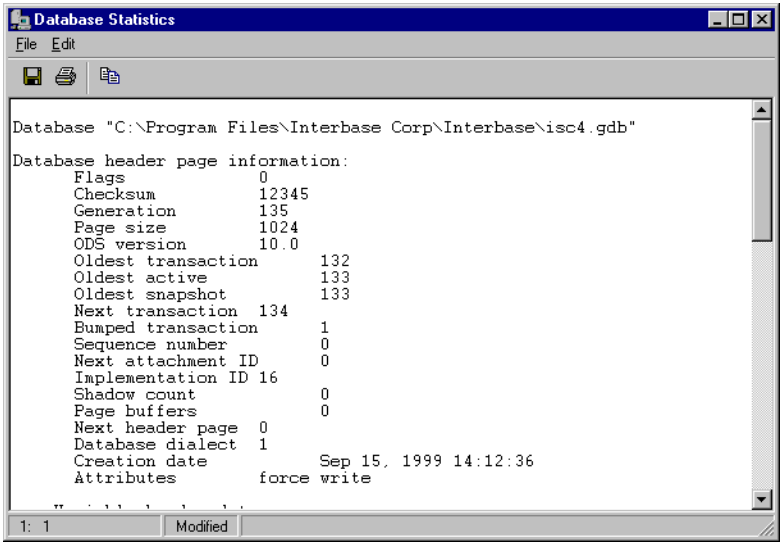

The Database Statistics report dialog is a standard text display window that exhibits database summary and database analysis information statistics. For an explanation of how to use the standard text display window, see "Standard text display window" on page 2-7.

Database statistics options

When you request a statistic option, InterBase generates and displays information for that database statistic. Possible statistic option values include: All Options, Data Pages, Database Log, Header Pages, Index Pages, and System Relations.

Note In addition to the selected statistic, header page information is displayed, regardless which statistic has been selected to report. If Header Pages is the selected option value, then only header page information will be displayed.

All Options

Displays statistic information for all options including Data Pages, Database Log, Header Pages, Index Pages, and System Relations.

This function corresponds to the **-all** option of **gstat**.

Data Pages

Displays data page information in the database summary. Below is an example of data page information, followed by an explanation of each item.

```
COUNTRY (31)
   Primary pointer page: 246, Index root page: 247
   Data pages: 1, data page slots: 1, average fill: 59%
```

```
Fill distribution:
0 - 19% = 0
20 - 39% = 0
40 - 59% = 1
60 - 79% = 0
80 - 99% = 0
```

The first line displays a database table name while the remaining lines contain item information pertaining to the table. These items include:

Table 9.2 Data page information

Item	Description
Primary pointer page	The page that is the first pointer page for the table.
Index root page	The page number that is the first pointer page for indexes.
Data pages	The total number of data pages.
Data page slots	The number of pointers to database pages, whether the pages are still in the database or not.
Average fill	The average percentage to which the data pages are filled.
Fill distribution	A histogram that shows the number of data pages that are filled to the percentages.

Database Log

Displays the database log in the database summary. Below is an example of database log information.

This function corresponds to the **-log** option of **gstat**.

```
Database log page information:
   Creation date Dec 20, 1998 11:38:19
   Log flags:2
      No write ahead log
   Next log page:0
   Variable log data:
   Control Point 1:
      File name:
      Partition offset: 0 Seqno: 0 Offset: 0
   Control Point 2:
      File name:
      Partition offset: 0 Seqno: 0 Offset: 0
   Current File:
      File name:
      Partition offset: 0 Seqno: 0 Offset: 0
```

Header Pages

Displays header page information in the database summary. Below is an example of database summary header page information, followed by an explanation of each item.

This function corresponds to the **-header** option of **gstat**.

```
Database "C:\Program Files\Borland\InterBase\examples\Database\employee.gdb"

Database header page information:
     Flags                0
     Checksum             12345
     Generation           41
     Page size            4096
     ODS version          11.0
     Oldest transaction   29
     Oldest active        30
     Oldest snapshot      30
     Next transaction     34
     Bumped transaction   1
     Sequence number      0
     Next attachment ID   0
     Implementation ID    16
     Shadow count         0
     Page buffers         0
     Next header page     0
     Database dialect     1
     Creation date        Aug 26, 2002 17:05:03

Variable header data:
     Sweep interval:      20000
     *END*

Service ended at 9/3/2002 4:59:05 PM
```

The first line displays the name and location of the primary database file while the remaining lines contain information on the database header page. These items include:

Table 9.3 Header page information

Item	Description
Checksum	The InterBase server no longer maintains a checksum values. The “12345” that appears in the header is bogus.
Generation	Counter incremented each time header page is written.
Page size	The current database page size, in bytes.
ODS version	The version of the database’s on-disk structure.

Table 9.3 Header page information

Item	Description
Oldest transaction	The transaction ID number of the oldest "interesting" transaction (those that are active, in limbo, or rolled back, but not committed).
Oldest active	The transaction ID number of the oldest active transaction.
Next transaction	The transaction ID number that InterBase assigns to the next transaction. The difference between the oldest transaction and the next transaction determines when database sweeping occurs. For example, if the difference is greater than this difference (set to 20,000 by default), then InterBase initiates a database sweep. See "**Overview of sweeping**" on page 6-21.
Sequence number	The sequence number of the header page (zero is used for the first page, one for second page, and so on).
Next connection ID	ID number of the next database connection.
Implementation ID:	The architecture of the system on which the database was created. These ID definitions are platform-dependent #define directives for a macro class named CLASS: • 1 HP Apollo Domain OS • 2 Sun Solaris SPARC, HP9000 s300, Xenix, Motorola IMP UNIX, UnixWare, NCR UNIX, NeXT, Data General DG-UX Intel • 3 Sun Solaris x86 • 4 VMS • 5 VAX Ultrix • 6 MIPS Ultrix • 7 HP9000 s700/s800 • 8 Novell NetWare • 9 Apple Macintosh 680x0 • 10 IBM AIX POWER series, IBM AIX PowerPC • 11 Data General DG-UX 88K • 12 HP MPE/xl • 13 SGI IRIX • 14 Cray • 15 SF/1 • 16 Microsoft Windows 32-bit Intel • 17 IBM OS/2 • 18 Microsoft Windows 16-bit • 19 Linux Intel • 20 Linux SPARC
Shadow count	The number of shadow files defined for the database.
Number of cache buffers	The number of page buffers in the database cache.

Table 9.3 Header page information

Item	Description
Next header page	The ID of the next header page.
Database dialect	The SQL dialect of the database
Creation date	The date when the database was created.
Attributes	• force write—indicates that forced database writes are enabled. • no_reserve—indicates that space is not reserved on each page for old generations of data. This enables data to be packed more closely on each page and therefore makes the database occupy less disk space. • shutdown—indicates database is shut down.
Variable header data	• sweep interval • secondary file information

Index Pages

Displays index information in the database summary. Below is an example of index page information, followed by an explanation of each item.

```
Index CUSTNAMEX (2)
   Depth: 2, leaf buckets: 2, nodes: 27
   Average data length: 45.00, total dup: 0, max dup: 0
   Fill distribution:
        0 - 19% = 0
       20 - 39% = 0
       40 - 59% = 1
       60 - 79% = 0
       80 - 99% = 1
```

Table 9.4 Index pages information

Item	Description
Index	The name of the index.
Depth	The number of levels in the index page tree. If the depth of the index page tree is greater than three, then sorting may not be as efficient as possible. To reduce the depth of the index page tree, increase the page size. If increasing the page size does not reduce the depth, then return it to its previous size.
Leaf buckets	The number of leaf (bottom level) pages in the index page tree.
Nodes	The total number of index pages in the tree.

Table 9.4 Index pages information

Item	Description
Average data length	The average length of each key, in bytes.
Total dup	The total number of rows that have duplicate indexes.
Max dup	The number of duplicates of the index with the most duplicates
Fill distribution	A histogram that shows the number of index pages filled to the specified percentages.

System Relations

Displays information for system tables in the database.

```
RDB$CHECK_CONSTRAINTS (24)
    Primary pointer page: 54, Index root page: 55
    Data pages: 5, data page slots: 5, average fill: 59%
    Fill distribution:
       0 - 19% = 0
      20 - 39% = 1
      40 - 59% = 0
      60 - 79% = 4
      80 - 99% = 0

   Index RDB$INDEX_14 (0)
   Depth: 1, leaf buckets: 1, nodes: 68
   Average data length: 0.00, total dup: 14, max dup: 1
   Fill distribution:
        0 - 19% = 0
       20 - 39% = 0
       40 - 59% = 1
       60 - 79% = 0
       80 - 99% = 0
```

The statistics contained here are similar to that of data pages and index pages. For information on the items see "Data Pages" and "Index Pages" above.

Monitoring client connections with IBConsole

You can view a list of users currently connected to a particular database in IBConsole using the Database Connections dialog. You can access this dialog by one of the following methods:

- Select a database (or any branch under the database hierarchy) in the Tree pane and choose **Database | Connected Users**.

- Select a database in the Tree pane and double-click **Connected Users** in the Actions column of the Work pane.
- Right-click a database in the Tree pane and choose **Connected Users** from the context menu.

Figure 9.3 Database connections dialog

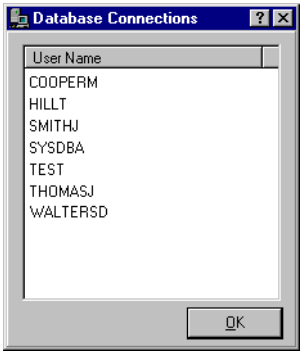

Note InterBase's temporary system tables provide resources for more extensive monitoring of database activity. See "Monitoring with system temporary tables" on page 9-1 of this chapter

The gstat command-line tool

Syntax `gstat [options] database`

Description The **gstat** program is a command-line tool for retrieving and reporting database statistics. Its function is the same as that described for IBConsole earlier in this chapter.

You must be SYSDBA or the owner of a database to run **gstat**. On UNIX platforms, there is a further constraint on **gstat**: In order to run **gstat**, you must have system-level read access to the database files. You can gain this by logging in as the same account that the InterBase server is running as (interbase or root) or by setting the system-level permissions on the database file to include read permission for your Group. These restrictions exist on UNIX platforms because **gstat** accesses the database file at the system level rather than through the InterBase server.

Note You can run **gstat** only against local databases: run **gstat** on the server host.

Options Table 9.5 lists the valid options for gstat.

Table 9.5 gstat options

Option	Description
-all	Equivalent to supplying **-index** and **-data**; this is the default if you supply none of **-index**, **-data**, or **-all**
-data	Retrieves and displays statistics on data tables in the database
-header	Stops reporting statistics after reporting the information on the header page
-index	Retrieves and displays statistics on indexes in the database
-log	Stops reporting statistics after reporting the information on the log pages
-pa[ssword] ***text***	Checks for password *text* before accessing a database
-r[ecord]	Used with -t, adds lines for average record length and average version length to the table statistics
-system	Retrieves statistics on system tables and indexes in addition to user tables and indexes
-t[able]	Outputs index and fill information for the requested table, in addition to database header, file, and log statistics
-user ***name***	Checks for user *name* before accessing database
-z	Prints product version of **gstat**

Example The following command requests table statistics, including record and version length for the JOB table in *employee.gdb*:

```
gstat -user SYSDBA -pa masterkey employee.gdb -t JOB -r
```

The command produces the following output:

```
Database "employee.gdb"

Database header page information:
        Flags                   0
        Checksum                12345
        Generation              26
        Page size               4096
        ODS version             11.0
        Oldest transaction      19
        Oldest active           20
        Oldest snapshot         20
        Next transaction        21
        Bumped transaction      1
        Sequence number         0
        Next attachment ID      0
```

```
        Implementation ID       16
        Shadow count            0
        Page buffers            0
        Next header page        0
        Database dialect        1
        Creation date           Sep 25, 2002 10:06:33

    Variable header data:
        Sweep interval:         20000
        *END*


Database file sequence:
File employee.gdb is the only file

Database log page information:
        Creation date
        Log flags:      2
                No write ahead log

        Next log page:  0

    Variable log data:
        Control Point 1:
                File name:
                Partition offset: 0     Seqno: 0        Offset: 0
        Control Point 2:
                File name:
                Partition offset: 0     Seqno: 0        Offset: 0
        Current File:
                File name:
                Partition offset: 0     Seqno: 0        Offset: 0
        *END*

Analyzing database pages ...

JOB (129)
    Primary pointer page: 178, Index root page: 179
    Average record length: 64.87, total records: 31
    Average version length: 0.00, total versions: 0, max versions: 0
    Data pages: 3, data page slots: 3, average fill: 72%
    Fill distribution:
         0 - 19% = 0
        20 - 39% = 1
        40 - 59% = 0
        60 - 79% = 0
        80 - 99% = 2

    Index MAXSALX (2)
```

```
        Depth: 1, leaf buckets: 1, nodes: 31
        Average data length: 4.00, total dup: 5, max dup: 1
        Fill distribution:
             0 - 19% = 1
            20 - 39% = 0
            40 - 59% = 0
            60 - 79% = 0
            80 - 99% = 0

    Index MINSALX (1)
        Depth: 1, leaf buckets: 1, nodes: 31
        Average data length: 4.00, total dup: 7, max dup: 2
        Fill distribution:
             0 - 19% = 1
            20 - 39% = 0
            40 - 59% = 0
            60 - 79% = 0
            80 - 99% = 0

    Index RDB$FOREIGN3 (3)
        Depth: 1, leaf buckets: 1, nodes: 31
        Average data length: 1.00, total dup: 24, max dup: 20
        Fill distribution:
             0 - 19% = 1
            20 - 39% = 0
            40 - 59% = 0
            60 - 79% = 0
            80 - 99% = 0

    Index RDB$PRIMARY2 (0)
        Depth: 1, leaf buckets: 1, nodes: 31
        Average data length: 10.00, total dup: 0, max dup: 0
        Fill distribution:
             0 - 19% = 1
            20 - 39% = 0
            40 - 59% = 0
            60 - 79% = 0
            80 - 99% = 0


[c:\program files\borland\interbase\examples\database]
```

Viewing lock statistics

Locking is a mechanism that InterBase uses to maintain the consistency of the database when it is accessed by multiple users. The lock manager is a thread in the **ibserver** process that coordinates locking.

The lock manager uses a lock table to coordinate resource sharing among client threads in the **ibserver** process connected to the database. The lock table contains information on all the locks in the system and their states. The global header information contains useful aggregate information such as the size of the lock table, the number of free locks, the number of deadlocks, and so on. There is also process information such as whether the lock has been granted or is waiting. This information is useful when trying to correct deadlock situations.

Syntax `iblockpr [a,o,w]` (Windows) or `gds_lock_print [a,o,w]` (UNIX)

`iblockpr [-i{a,o,w}] [t n]`

Description **iblockpr** monitors performance by checking lock requests.

The first form of syntax given above retrieves a report of lock statistics at one instant in time. The second form monitors performance by collecting samples at fixed intervals.

The options display interactive information on current activity in the lock table. The utility prints out the events per second for each sampling and gives the average of the values in each column at the end.

Table 9.6 iblockpr/gds_lock_print options

Option	Description
[**none**]	Same as ***-o***
-a	Prints a static view of the contents of the lock table
-o	Prints a static lock table summary and a list of all entities that own blocks
-w	Prints out all the information provided by the **-o** flag plus wait statistics for each owner; this option helps to discover which owner's request is blocking others in the lock table
The following options supply interactive statistics (events/second) for the requested items, which are sampled *n* times every *t* seconds, with one line printed for each sample. The average of the sample values is printed at the end of each column. If you do not supply values for *n* and *t*, the default is *n*=1.	
-i	Prints all statistics; the output is easier to read if you issue **-ia**, **-io**, and **-iw** separately
-ia	Prints how many threads are trying to acquire access to the lock table per second
-io	Prints operation statistics such lock requests, conversions, downgrades, and releases per second

Table 9.6 iblockpr/gds_lock_print options

Option	Description
-iw	Prints number of lock acquisitions and requests waiting per second, wait percent, and retries
t	Specifies the time in seconds between samplings
n	Specifies the number of samples to be taken

Example The following statement prints "acquire" statistics (access to lock table: acquire/s, acqwait/s, %acqwait, acqrtry/s, and rtrysuc/s) every three seconds until ten samples have been taken:

```
gds_lock_print -ia 3 10
```

Retrieving statistics with *api_database_info*()

InterBase includes programming facilities to gather performance timings and database operation statistics.

You can use the API function *isc_database_info( )* to retrieve statistics, by specifying one or more of the following request buffer items:

Table 9.7 Database I/O statistics information items

Request Buffer Item	Result Buffer Contents
isc_info_fetches	Number of reads from the memory buffer cache; calculated since the InterBase server started
isc_info_marks	Number of writes to the memory buffer cache; calculated since the InterBase server started
isc_info_reads	Number of page reads; calculated since the InterBase server started
isc_info_writes	Number of page writes; calculated since the InterBase server started
isc_info_backout_count	Number of removals of record versions per table since the current database attachment started
isc_info_delete_count	Number of row deletions • Reported per table • Calculated since the current database attachment started

Table 9.7 Database I/O statistics information items (*continued*)

Request Buffer Item	Result Buffer Contents
isc_info_expunge_count	Number of removals of a record and all of its ancestors, for records whose deletions have been committed • Reported per table • Calculated since the current database attachment started
isc_info_insert_count	Number of inserts into the database • Reported per table • Calculated since the current database attachment started
isc_info_purge_count	Number of removals of old versions of fully mature records (records committed, resulting in older-ancestor-versions no longer being needed) • Reported per table • Calculated since the current database attachment started
isc_info_read_idx_count	Number of reads done via an index • Reported per table • Calculated since the current database attachment started
isc_info_read_seq_count	Number of sequential database reads, that is, the number of sequential table scans (row reads) • Reported per table • Calculated since the current database attachment started
isc_info_read_update_count	Number of row updates • Reported per table • Calculated since the current database attachment started

See the *API Guide* for information on request buffers, and details of using this API call.

Chapter

10

Interactive Query

This chapter documents the IBConsole interactive SQL (ISQL) and command-line **isql** utilities for InterBase. These tools provide an interface to InterBase's Dynamic SQL interpreter. You can use these query tools to perform data definition, prototype queries before implementing them in your application, or to perform ad hoc examination of data in your database.

Topics covered in this chapter include:

- IBConsole ISQL
- Executing SQL statements
- Committing and rolling back transactions
- Saving ISQL input and output
- Changing ISQL settings
- Extracting metadata
- Command-line isql tool

IBConsole ISQL

The IBConsole ISQL Window permits you to execute DDL and DML commands to the InterBase server as well as to load, save, print, cut, paste, and copy SQL scripts and results.

The interactive SQL window

To display the ISQL Window:

- Click the Interactive SQL toolbar button ().
- Choose **Tools | Interactive SQL**.

Figure 10.1 The interactive SQL editor in IBConsole

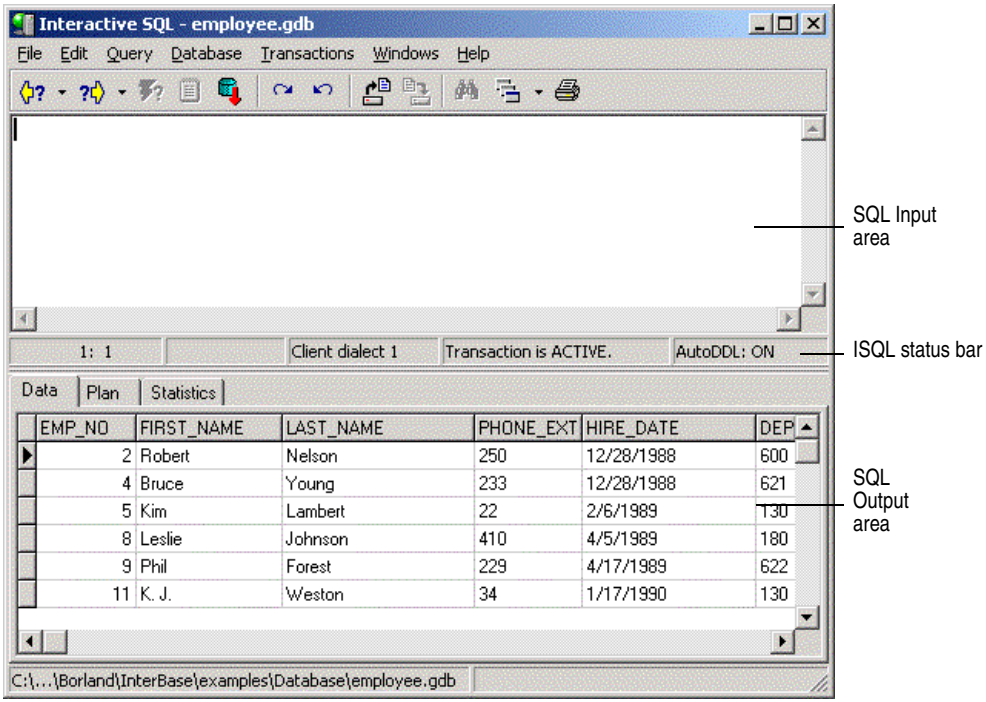

ISQL menus

The IBConsole Interactive SQL window—also called the SQL Editor—are the File, Edit, Query, Database, Transactions, and Windows menus.

File menu

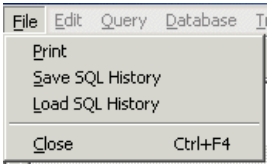

The File menu contains the following entries:

- **Print** prints the current contents of the SQL Input area.
- **Save** saves the stack of SQL commands from the current session to a file.
- **Load** loads a saved SQL history file.
- **Close** closes the SQL Editor.

Edit menu

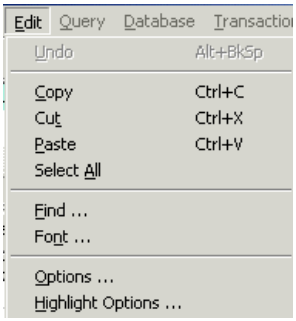

The Edit menu contains the following entries:

- **Undo**, **Cut**, and **Paste** are for use in the SQL Input area only. “Undo” in the Edit menu does not undo database changes. Use **Transactions | Rollback** to undo database changes.
- **Copy**, **Select All**, and **Find** are for use in both the Input and Output area.
- **Font** specifies the font to be used in the Input area
- **Options** are described in Table 10.2 on page 10-12. They include settings for query plan, auto commit, Blobs, and terminators.
- **Highlight Options** allows you to set the appearance of text in the Input area

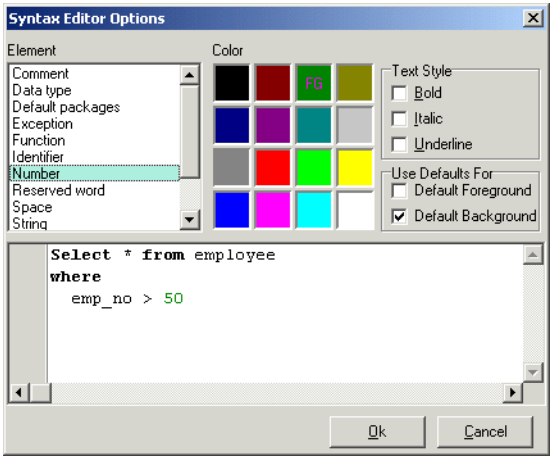

Query menu

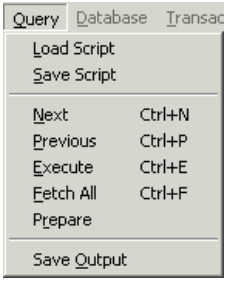

The Query menu contains the following entries:

- **Load Script** loads a SQL script file
- **Save Script** saves a SQL script file
- **Next** steps you forward to the next SQL statement after you have used the Previous command to display earlier statements in the stack.
- **Previous** displays the previous statement that you entered in the Input area. This command can be used repeatedly to step through the stack of statements entered during the current interactive SQL session.
- **Execute** executes the SQL statements currently displayed in the Input area.
- **Fetch All** forces the IBConsole to fetch all the rows of a subquery rather than a subset. Rarely needed, because IBConsole fetches the results as you scroll down in the Output area.
- **Prepare** displays the optimization plan in the SQL output area.
- **Save Output** saves the contents of the Output area to a text file.

Database menu

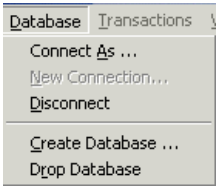

The Database menu contains the following entries:

- **Connect As** allows you to connect to the database as a different user.
- **New Connection** is active only when you are not connected to a database. It displays the Database Connect dialog.

- **Create Database** and **Drop Database** are exactly the same as the equivalent commands in the main IBConsole **Database** menu. They are provided in the SQL Editor **Database** menu for convenience.

Transactions menu

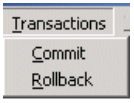

The Transactions menu allows you to commit or roll back transactions.

Windows menu

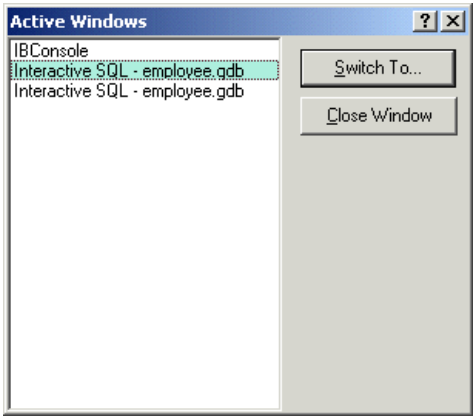

Clicking the Windows menu displays a list of current open windows. You can display a window by highlighting it in the list and clicking Switch to or by double-clicking the entry.

ISQL toolbar

The IBConsole SQL Editor toolbar contains the following buttons:

Table 10.1 Toolbar buttons for executing SQL statements

Button	Description
	Displays the previous statement that you entered in the Input area. This button can be used repeatedly to step through the stack of statements entered during the current interactive SQL session. The accelerator key is Ctrl-P. Click the down arrow next to the button to see a list of available statements. Click a statement in the list to display it in the Input area: 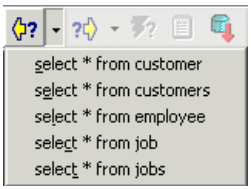 (Same as **Query** | **Previous**.)
	Steps you forward to the next SQL statement after you have used the Previous command to display earlier statements in the stack. The accelerator key is Ctrl-N. Click the down arrow next to the button to see a list of available statements. Click a statement in the list to display it in the Input area. (Same as **Query** | **Next**.)
	Executes the current statement or script in the SQL Input area. The output is displayed in the SQL Output area. The accelerator key is Ctrl-E. (Same as **Query** | **Execute**.)
	Displays the query plan for the current query in the SQL output area. (Same as **Query** | **Prepare**.)
	Forces IBConsole to fetch all the rows of a subquery rather than a subset. Rarely needed, because IBConsole fetches the results as you scroll down in the Output area. (Same as **Query** | **Fetch All**.)
	Commits the transaction specified by the SQL statement to the database. (Same as **Transactions** | **Commit**.)
	Rolls back all database changes since the last commit. (Same as **Transactions** | **Commit**.)
	Loads a script for SQL execution into the SQL input area. (Same as **Query** | **Load Script**.)
	Saves SQL statements entered in the SQL input area to a file. (Same as **Query** | **Save Script**.)

Table 10.1 Toolbar buttons for executing SQL statements

Button	Description
	Finds text in the SQL Input area. (Same as **Edit** | **Find**.)
	Displays a list of current open windows. You can display a window by highlighting it in the list and clicking Switch To or by double-clicking the entry. You can also change windows by clicking the down arrow and clicking the window you want from the dropdown list: 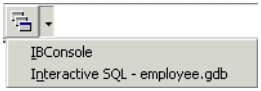
	Prints the contents of the SQL Input area.

ISQL work areas

SQL Input area

The SQL input area is where you can type SQL statements or scripts to be executed. It scrolls vertically.

SQL Output area

The SQL output area is where the results of the SQL statements or scripts are displayed. It scrolls both horizontally and vertically. The SQL output area contains two tabs:

- The Data tab displays any data returned by the SQL output in a grid format.
- The Plan tab displays the plan for the SQL statement or script.
- The Statistics tab displays statistics for the SQL output, including: execution time, prepare time, starting memory, delta memory, current memory, number of buffers, number of reads, number of writes, the plan, and the number of records fetched.

Status bar

The status bar at the bottom of the SQL input area displays information relevant to the SQL input areas such as cursor position, input status, client dialect, and transaction status. You can change the client dialect by right clicking on the status bar.

Managing ISQL temporary files

ISQL creates temporary files used during a session to store information such as the command history, output file names, and so on. These files are named in the form *isql_aa.xx*. The files are stored in the directory specified by the TMP environment variable, or if that is not defined, the working directory, or if that is not defined, they are stored in the *Windows* directory.

To avoid cluttering the *Windows* directory with InterBase temporary files, specify a different directory for them by defining TMP.

When you exit, ISQL deletes these temporary files. If ISQL abnormally terminates (for example, due to a power failure), then these files remain and may be freely deleted without any adverse effects. You should not delete any of these temporary files while ISQL is running, because they may be used in the current session.

Executing SQL statements

In ISQL, you can execute SQL statements in either of two ways:

- Interactively, one statement at a time
- From a script containing multiple statements

Executing SQL interactively

To execute a SQL statement interactively:

1 Type a single SQL statement in the SQL Input area. Make sure any other existing statements are *commented*. A statement is commented if it is preceded by "/*" and followed by "*/".

   If the statement already exists in the SQL Input area make sure all statements except the one you wish to execute are commented. Commented statements in the SQL Input area are ignored during execution.

2 Choose **Query | Execute**, enter Ctrl+E, or click the Execute toolbar button.

   If more than one statement is uncommented, Execute executes each statement, one after the other.

Tip You can copy text from other Windows applications such as the Notepad and Wordpad text editors and paste it into the SQL Input area. You can also copy statements from the ISQL Output area and paste them into the SQL Input area. This cut-and-paste method is also a convenient way to use the online SQL tutorial provided in the online Help.

When SQL statements are executed, whether successfully or not, they become part of the ISQL *command history*, a sequential list of SQL statements entered in the current session.

Preparing SQL statements

Use the Prepare toolbar button, or select **Query | Prepare**, to prepare SQL statements for execution and to view the query plan. Prepare compiles the query plan on the server, and displays it in the Plan tab of the SQL output area. Use Prepare to determine if your SQL script is well-constructed, without having to wait for the SQL script to execute.

Valid SQL statements

- You can execute interactively any SQL statement identified as "available in DSQL" in the *Language Reference*. You cannot use any statements that are specifically identified in the *Language Reference* as **isql** statements; all these have functionally equivalent menu items in ISQL.

  For example, the SET NAMES statement cannot be executed from the SQL Input area. To change the active character set, choose **Edit | Options** and select the desired character set option value in the SQL Options dialog.

- SQL script files can include statements that are not valid to enter interactively. For example, you can use the SET statements such as SET LIST or SET TERM in scripts.

- Transaction names may not be used with SET TRANSACTION statement.

- The SQL Input area accepts multiple statements, although only one can be executed at a time. Each statement entered in the SQL input area must be terminated by a semicolon (;). The SQL Input area accepts multiple statements, although only one can be executed at a time. An uncommented statement that holds the mouse cursor is called the *current statement*.

Loading and executing a SQL script file

To execute a SQL script file containing SQL statements:

1. Choose **Query | Load Script** or click the Load Script toolbar button.
2. Locate the desired script file in the Open dialog, and click Open to display the statements of the script file in the SQL Input area.
3. Ensure that you are connected to the desired database.
4. If you are connected to the database, comment out any CONNECT or CREATE DATABASE statements.
5. Choose **Query | Execute** or click **Execute** on the toolbar to begin executing the entire script statement by statement.

Note Statements executed from a loaded script file do not become part of the command history.

Committing and rolling back transactions

Changes to the database from data definition (DDL) statements—for example, CREATE and ALTER statements—are automatically committed by default. To turn off automatic commit of DDL, choose **Edit | Options** and set the Auto Commit DDL option to false in the SQL Options dialog.

Changes made to the database by data manipulation (DML) statements—for example INSERT and UPDATE—are not permanent until they are committed. Commit changes by choosing **Transactions | Commit** or by clicking Commit on the toolbar.

To undo all database changes from DML statements since the last commit, choose **Transactions | Rollback** or click Rollback on the toolbar.

Saving ISQL input and output

You can save the following to a file:

- SQL statements entered in the SQL Input area of the current session.
- The output of the last SQL statement executed.

Saving SQL input

To save the SQL statements entered in the SQL Input area of the current session to a text file:

1. In the SQL Editor, choose **Query | Save Script** or click the Save Script toolbar button.
2. Enter a file name, including the location for the new file, in the Save As dialog and click Save.

   To include the location for the file, type the file path and file name in the Filename text area, or browse to the folder where you would like the file to reside and type only the file name.

Only the SQL statements entered in the current session, not the output, are saved to the specified file.

Saving SQL output

To save the results of the last executed SQL statement to a file:

1. In the SQL Editor, choose **Query | Save Output**.
2. Enter a file name, including the location for the new file, in the Export To dialog and click Save.

   To include the location for the file, either type the file path and file name in the Filename text area, or browse to the folder where you would like the file to reside and type only the file name.

The output in the Data tab from the last successful statement is saved to the named text file.

If you run a SQL script, and then choose to save the output, all the commands in the script file and their results are saved to the output file. If command display has been turned off in a script with SET ECHO OFF, then SQL statements in the script are not saved to the file.

Changing ISQL settings

Use the SQL Options dialog to display and modify ISQL session settings, determine if the main IBConsole window will be updated based on the statements given in the ISQL window, and specify how open transactions are handled when the ISQL window is closed.

Select **Edit | Options** from the Interactive SQL window to display the SQL Options dialog.

The SQL Options dialog has two tabs: Options and Advanced.

Options tab

Use the Options tab to display and modify the ISQL session settings. You can specify options by clicking the option value and choosing a new value from a drop down list of values or by double-clicking the option value to rotate its value to the next in the list of values.

Figure 10.2 Options tab of the SQL Options dialog

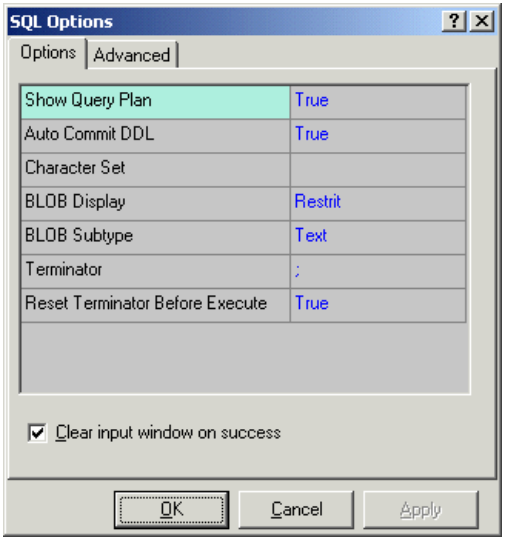

The following table summarizes the **isql** session settings available on the Options tab.

Table 10.2 Options Tab of the SQL Options dialog

Setting	Description
Show Query Plan	Values: true (default) or false If this setting is True, IBConsole displays the query plan chosen by the optimizer when a SELECT statement is entered. To modify the optimizer plan, use the PLAN option of the SQL SELECT statement. See "SET PLAN" on page 10-40.
Auto Commit DDL	Values: true (default) or false • If this setting is True, IBConsole automatically commits DDL (data definition) statements as each statement is entered. • If this setting is False, you must explicitly commit DDL statements (with **Transactions | Commit**) to make them permanent. See "SET AUTODDL"on page 10-34.
Character Set	Determines the active character set for strings for subsequent connections to the database; enables you to override the default character set for a database. • Specify the character set before connecting to the database whose character set you want to specify. For a complete list of character sets recognized by InterBase, see the *Language Reference*. • Your choice of character set limits possible collation orders to a subset of all available collation orders. Given a character set, a collation order can be specified when data is selected, inserted, or updated in a column. • You can perform the same function in a SQL script with the SET NAMES command. Use SET NAMES before connecting to the database whose character set you want to specify. See "SET NAMES" on page 10-40 for more information.
BLOB Display	• Values: Enabled (default), Disabled, Restrict Determines how IBConsole displays columns of Blob data. SELECT always displays the Blob ID for columns of Blob datatype. By default, a SELECT also displays actual Blob data of text subtypes beneath the associated row. • If this setting is set to Enabled, IBConsole displays the contents of Blob columns. • If this setting is set to Disabled, IBConsole does not display the contents of Blob columns. If this setting is set to Restrict, IBConsole displays the contents of only Blob columns of the specified BLOB Subtype.
Terminator	Identifies the end-of-statement symbol to be used for SQL queries

Table 10.2 Options Tab of the SQL Options dialog (*continued*)

Setting	Description
Clear input window on success	Check this box to clear the SQL input window after a SQL statement is successfully executed

Advanced tab

Use the Advanced tab to determine if the main IBConsole window will be updated based on the statements given in the ISQL window, and to specify how open transactions are handled when the ISQL window is closed.

Figure 10.3 Advanced tab of the SQL Options dialog

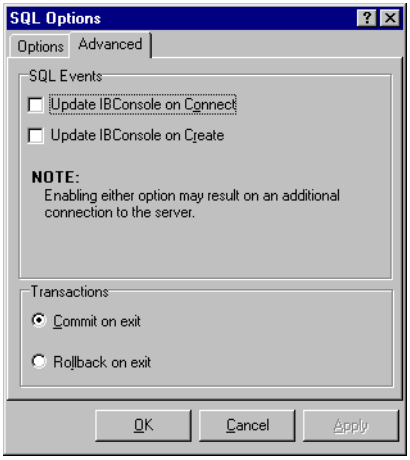

The following table summarizes the settings available on the Advanced tab:

Table 10.3 Advanced tab of the SQL Options dialog

Setting	Description
Update IBConsole on Connect	Does not update the connected databases in the IBConsole window; however, it ensures that the IBConsole window is updated to reflect statements executed in the ISQL window.

Table 10.3 Advanced tab of the SQL Options dialog

Setting	Description
Update IBConsole on Create	Updates the main window if the currently selected server is active. Automatically registers any database created in the ISQL window and creates an alias for it.
Commit on exit	Commits any active transactions when the ISQL window is closed.
Rollback on exit	Rolls back any active transactions when the ISQL window is closed.

Inspecting database objects

Use the object inspector to view properties, metadata, permissions, data, and dependencies for the entire database or for a specific table, view, function, procedure, or any other database attribute displayed in the Tree pane.

To open the object inspector, double-click a database object in the Work pane. The object inspector appears:

Figure 10.4 Object inspector

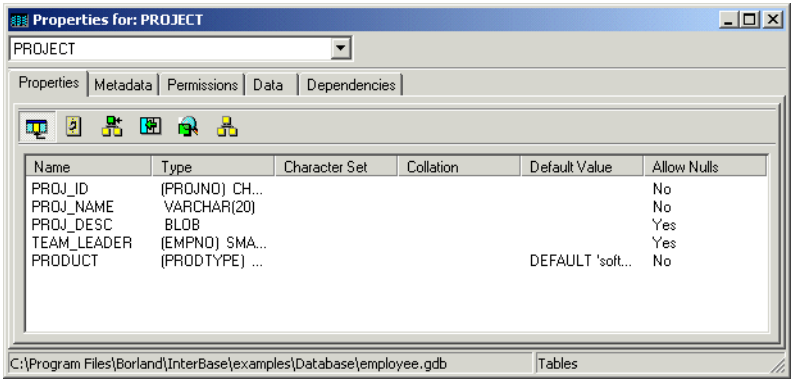

Depending on the database object selected, the object inspector has some or all of the following tabs: Properties, Metadata, Permissions, Data, and Dependencies. These are discussed in the following sections.

Viewing object properties

The Properties tab is available when viewing Table and View database objects. Use the Properties tab of the object inspector to display properties for database objects, including columns, triggers, check constraints, indexes, unique constraints, and referential constraints. The Properties tab has five toolbar buttons for displaying the various object properties:

Table 10.4 Object inspector toolbar buttons

Button	Description
	**Show columns**: displays the name, type, collation, character set, default value, and whether or not null values are acceptable for every row in the column. The accelerator key is Ctrl+Alt+C. For more information on columns, refer to "Defining columns" in the *Data Definition Guide*.
	**Show triggers**: displays the name and type of each trigger, as well as whether or not it is active. In addition, it displays the SQL trigger statement. The accelerator key is Ctrl+Alt+T. For more information on triggers, refer to "Working with Triggers" in the *Data Definition Guide*.
	**Show check constraints**: displays the names of the constraints, whether or not they can be deferred, and if they were initially deferred. In addition, it displays the SQL check constraint statements. The accelerator key is Ctrl+Alt+H. For more information, refer to "Defining a CHECK constraint" in the *Data Definition Guide*.
	**Show indexes**: displays the name of the index keys, and whether or not they are unique, descending, or active. The accelerator key is Ctrl+Alt+R. For more information, refer to "Working with Indexes" in the *Data Definition Guide*.
	**Show unique constraints**: displays the names of the constraints, whether or not they can be deferred, if they were initially deferred, and the index keys. The accelerator key is Ctrl+Alt+U.
	**Show referential constraints**: displays the names of the constraints, whether or not they can be deferred, if they were initially deferred, the update rule, the update rules, the delete rules, the index, and the reference table. The accelerator key is Ctrl+Alt+R.

Viewing metadata

The metadata which the Metadata tab of the object inspector displays depends on the database that is selected in the Tree pane, or the item that is selected in the Work pane.

To view metadata for an entire database Select a connected database in the Tree pane, and then double-click View Metadata in the Work pane. The metadata is displayed in a text window.

To view metadata for a specific database object perform one of the following actions:

- Select a database element from the hierarchy displayed in the Tree pane, and then in the Work pane double-click an object to display its Properties dialog. Click the Metadata tab to see the object's metadata.
- Select a database element from the hierarchy displayed in the Tree pane, and then in the Work pane right-click a database object associated with that element and select Extract from the context menu.

For example, if you want metadata for domains only, expand the desired database hierarchy (if it is not already expanded), select Domains, double-click on a domain in the Work pane, and select the Metadata tab of the Properties dialog.

Use the drop down list at the top of the dialog to select other objects associated with the database element.

The following table lists the items for which you can view metadata for associated objects with the object inspector.

Table 10.5 Metadata information items

Item	Displays
Blob Filters	Blob filters definition
Domains	Metadata script, dependencies, datatype, description, check constraints, and default values
Exceptions	Description, exception number, exception message, metadata script, and dependencies
External Functions	UDFs definition
Generators	Generator ID, current value, metadata script, and dependencies
Stored Procedures	Metadata script, procedure body, input parameters, output parameters, permissions, data, and dependencies
Roles	Role definition
Tables	Columns, datatypes, triggers, indexes, unique constraints, referential constraints, check constraints, metadata script, permissions, data, and dependencies
Views	Metadata script, permissions, data, and dependencies

Extracting metadata

You can extract a metadata script to a file by displaying the desired metadata in the Metadata tab and clicking the Save Script toolbar button.

Extracting an entire database exports metadata in a specific order, to allow the resulting script to be used as input to recreate the database.

Table 10.6 Metadata extraction constraints

Metadata	Comments
1. Database	Extracts database with default character set and PAGE_SIZE
2. Domains	Must be before tables that reference domains
3. Tables	Must be after domains
4. Indexes	Must be after tables
5. FOREIGN KEY constraints	Must be added after tables to avoid tables being referenced before they have been created
6. Views	Must be after tables
7. CHECK constraints	Must be after tables
8. Exceptions	Must be extracted before stored procedures and triggers that contain code to raise exceptions
9. Stored procedures	Stored procedures are shown with no body in CREATE PROCEDURE and then ALTER PROCEDURE to add the text of the procedure body; this is to allow circular or recursive procedure references
10. Triggers	Must be after tables Must be after stored procedures, to allow trigger code to reference procedures Does not extract triggers from CHECK constraints
11. Roles	Must be before GRANT privileges
12. GRANTS	Must be after tables, views, stored procedures, triggers, and roles

Items that are not extracted include:

- Code of external functions or filters, because that code is not part of the database. The declarations to the database (with DECLARE EXTERNAL FUNCTION and DECLARE FILTER) are extracted
- System tables, system views, and system triggers
- Because DDL statements do not contain references to object ownership, the extracted file does not show ownership. The output file includes the name of the object and the owner if one is defined. There is no way to assign an object to its original owner.

Extracting metadata

You can extract a metadata script to a file by displaying the desired metadata in the Metadata tab and clicking the Save Script toolbar button.

Extracting an entire database exports metadata in a specific order, to allow the resulting script to be used as input to recreate the database.

Table 10.7 Order of metadata extraction

Metadata	Comments
1. Database	Extracts database with default character set and PAGE_SIZE
2. Domains	Must be before tables that reference domains
3. Tables	Must be after domains
4. Indexes	Must be after tables
5. FOREIGN KEY constraints	Must be added after tables to avoid tables being referenced before they have been created
6. Views	Must be after tables
7. CHECK constraints	Must be after tables
8. Exceptions	Must be extracted before stored procedures and triggers that contain code to raise exceptions
9. Stored procedures	Stored procedures are shown with no body in CREATE PROCEDURE and then ALTER PROCEDURE to add the text of the procedure body; this is to allow circular or recursive procedure references
10. Triggers	Must be after tables Must be after stored procedures, to allow trigger code to reference procedures Does not extract triggers from CHECK constraints.
11. Roles	Must be before GRANT privileges
12. GRANTs	Must be after tables, views, stored procedures, triggers, and roles

Items that are not extracted include:

- Code of external functions or filters, because that code is not part of the database. The declarations to the database (with DECLARE EXTERNAL FUNCTION and DECLARE FILTER) are extracted.
- System tables, system views, and system triggers.

- Because DDL statements do not contain references to object ownership, the extracted file does not show ownership. The output file includes the name of the object and the owner if one is defined. There is no way to assign an object to its original owner.

Command-line isql tool

Command-line **isql** is a utility for processing SQL data definition (DDL) and data manipulation (DML) statements from interactive input or from a source file. It enables you to create and view metadata, add and modify data, grant user permissions, test queries, and perform database administration tasks.

This section provides an introduction to using **isql**. For a description of the standard SQL commands available in **isql**, see the *Language Reference*. For a description of special **isql** commands, see "isql command reference" on page 10-27.

Invoking isql

You can use **isql** in the following ways:

- Interactively to process SQL statements, by entering statements at the **isql** prompt
- Noninteractively to process SQL statements in a file

To start the isql utility, type the following at a UNIX shell prompt or Windows console prompt:

```
isql [options] [database_name]
```

where **options** are command-line options and *database_name* is the name of the database to connect to, including disk and directory path.

If no options are specified, **isql** starts an interactive session. If no database is specified, you must connect to an existing database or create a new one. If a database was specified, **isql** starts the interactive session by connecting to the named database.

If options are specified, **isql** starts interactively or noninteractively, depending on the options. For example, reading an input file and writing to an output file are noninteractive tasks, so the **-input** or **-output** options do not start an interactive session. Additional noninteractive options include **-a**, **-database**, **-extract**, and **-x**, which are used when extracting DDL statements.

When you start an interactive **isql** session, the following prompt appears:

```
SQL>
```

You must then end each command with a terminator character. The default terminator is a semicolon (;). You can change the terminator to any character or group of characters with the SET TERMINATOR command or with the **-terminator** command-line option. If you omit the terminator, a continuation prompt appears (CON>).

Note For clarity, all of the commands and examples in this chapter end with the default semicolon terminator.

Command-line options

Only the initial characters in an option are required. You can also type any portion of the text enclosed in brackets, including the full option name. For example, specifying **-n**, **-no**, or **-noauto** has the same effect.

Table 10.8 isql command-line options

Option	Description
-a	Extracts all DDL for the named database
-c[ache]	Set number of cache buffers for this connection to the database; see "Default cache size per ISQL connection" on page 6-25.
-d[atabase] *name*	Used with **-x**; changes the CREATE DATABASE statement that is extracted to a file • Without **-d**, CREATE DATABASE appears as a C-style comment and uses the database name specified on the **isql** command line • With **-d**, **isql** extracts an uncommented CREATE DATABASE and substitutes *name* as its database argument
-e[cho]	Displays (echoes) each statement before executing it
-ex[tract]	Same as **-x**
-i[nput] *file*	Reads commands from an input file instead of from standard input • Input files can contain **-input** commands that call other files, enabling execution to branch and then return • **isql** exits (with a commit) when it reaches the end of the first file • In interactive sessions, use **-input** to read commands from a file
-m[erge_stderr]	• Merges stderr output with stdout • Useful for capturing output and errors to a single file when running **isql** in a shell script or batch file
-n[oauto]	Turns off automatic commitment of DDL statements; by default, DDL statements are committed automatically in a separate transaction
-nowarnings	Displays warning messages if and only if an error occurs (be default, **isql** displays any message returned in a status vector, even if no error occurred)

Table 10.8 isql command-line options (*continued*)

Option	Description
-o[utput] file	Writes results to an output file instead of to standard output; in interactive sessions, use **-output** to write results to a file
-pas[sword] ***password***	Used with **-user** • Specifies a password when connecting to a remote server • For access, both *password* and *user* must represent a valid entry in the security database
-page[length] ***n***	Prints column headers every *n* lines instead of the default 20
-q[uiet]	
-r[ole] ***rolename***	Grants privileges of role *rolename* to *user* on connection to the database
-s[qldialect] ***n***	Interprets subsequent commands as dialect *n* until end of session or until dialect is changed by a SET SQL DIALECT statement • For *n* = 1, commands are processed as in InterBase 5 or earlier • For *n* = 2, elements that have different interpretations in dialect 1 and 3 are all flagged with warnings or errors to assist in migrating databases to dialect 3 • For *n* = 3, all statements are parsed as current InterBaseSQL semantics: double quotes are delimited identifiers, DATE datatype is SQL DATE, and exact numerics with precision greater than 9 are stored as INT64
-t[erminator] ***x***	Changes the end-of-statement symbol from the default semicolon (**;**) to *x*, where *x* is a single character or any sequence of characters; deprecated in InterBase 7
-u[ser] ***user***	Used with **-password**; specifies a user name when connecting to a remote server • For access, both *password* and *user* must represent a valid entry in the security database
-x	Extracts DDL for the named database; displays DDL to the screen unless redirected to a file
-z	Displays the software version of **isql**

Using warnings

Warnings can be issued for the following conditions:

- SQL statements with no effect
- SQL expressions that produce different results in InterBase 5 versus InterBase 6 or later
- API calls that will be replaced in future versions of the product

- Pending database shutdown

Examples

- Suppose *createdb.sql* contains DDL statements to create a database. To execute the statements, enter:

```
isql -input createdb.sql
```

- The following example starts an interactive connection to a remote database. The remote server, jupiter, accepts the specified user and password combination with the privileges assigned to the STAFF role:

```
isql -user sales -password mycode -role 'staff'
   'jupiter:/usr/customer.ib'
```

- The next example starts an interactive session but does not attach to a database. **isql** commands are displayed, and query results print column headers every 30 lines:

```
isql -echo -page 30
```

Exiting isql

To exit **isql** and roll back all uncommitted work, enter:

```
QUIT;
```

To exit **isql** and commit all work, enter:

```
EXIT;
```

Connecting to a database

If you do not specify a database on the command-line when invoking **isql**, you must either connect to an existing database or create a new one. Use the CONNECT command to connect to a database and CREATE DATABASE to create a database. For the full syntax of CONNECT and CREATE DATABASE, see the *Language Reference*.

You can connect to a database in two ways. You can connect to:

- A local database on a Windows platform. Use the CONNECT command with the full path of the database as the argument. For example:

```
SQL> CONNECT 'C:/Borland/InterBase/Database/examples/employee.gdb' role 'staff';
```

- A remote database on a Windows or UNIX server using TCP/IP. Use the CONNECT command with the full node name and path of the database as the argument. Separate the node name from the database path with a colon.

To connect to a database on a UNIX platform named jupiter:

```
SQL> CONNECT 'jupiter:/usr/interbase/examples/employee.gdb';
```

To connect to a database on a Windows 2000 platform named venus:

```
SQL> CONNECT 'venus:c:/InterBase/InterBase
   /examples/database/employee.gdb';
```

Note Be careful not to confuse node names and shared disks, since both are specified with a colon separator. If you specify a single letter that maps to a disk drive, it is assumed to be a drive, not a node name.

Tip You can use either forward slashes (/) or backslashes (\ ) as directory separators. InterBase automatically converts either type of slash to the appropriate type for the server operating system.

Setting isql client dialect

To use **isql** to create a database in a particular dialect, first set **isql** to the desired dialect and then create the database. You can set **isql** dialect the following ways:

- On the command line, start **isql** with option **-sql_dialect** *n*, where *n* is 1, 2, or 3:

```
isql -sql_dialect n
```

- Within an **isql** session or in a SQL script, include the following statement:

```
SET SQL DIALECT n;
```

isql dialect precedence is as follows:

- Lowest: Dialect of an attached version 6 or later database
- Next lowest: Dialect specified on the command line
- Next highest: Dialect specified during the session
- Highest: Dialect of an attached Version 5 database (=1)

In InterBase, **isql** has the following behavior with respect to dialects:

- If you start **isql** and attach to a database without specifying a dialect, **isql** takes on the dialect of the database.
- If you specify a dialect on the command line when you invoke **isql**, it retains that dialect after connection unless explicitly changed.
- When you change the dialect during a session using SET SQL DIALECT *n*, **isql** continues to operate in that dialect until it is explicitly changed.
- When you create a database using **isql**, the database is created with the dialect of the **isql** client; for example, if **isql** has been set to dialect 1, when you create a database, it is a dialect 1 database.
- If you create a database without first specifying a dialect for the **isql** client or attaching to a database, **isql** creates the database in dialect 3.

The statements above are true whether you are running **isql** as a command-line utility or accessing it through IBConsole.

Important Any InterBase **isql** client that attaches to a Version 5 database resets to dialect 1.

Transaction behavior in isql

When you start **isql**, InterBase begins a transaction. That transaction remains in effect until you issue a COMMIT or ROLLBACK statement. You must issue a COMMIT or ROLLBACK statement to end a transaction. Issuing one of these statements automatically starts a new transaction. You can also start a transaction with the SET TRANSACTION statement.

isql uses a separate transaction for DDL statements. When these statements are issued at the SQL> prompt, they are committed automatically as soon as they are completed. DDL scripts should issue a COMMIT after every CREATE statement to ensure that new database objects are available to all subsequent statements that depend on them. For more information on DDL statements, see the *Data Definition Guide*.

Extracting metadata

You can extract the DDL statements that define the metadata for a database to an output file with the **-extract** option. Adding the optional **-output** flag reroutes output to a named file. Use this syntax:

```
isql [[-extract | -x][-a] [[-output | -o] outputfile]] database;
```

The **-x** option is an abbreviation for **-extract**. The **-a** flag directs **isql** to extract all database objects. Note that the output file specification, *outputfile*, must follow the **-output** flag, while you can place the name of the database being extracted at the end of the command.

Table 10.9 isql extracting metadata arguments

Option	Description
database	File specification of the database from which metadata is being extracted
outputfile	File specification of the text file to receive the extracted statements; if omitted, **isql** writes the information to the screen

You can use the resulting text file to:

- Examine the current state of a database's system tables before you plan alterations to it, or when a database has changed significantly since its creation.
- Use your text editor to make changes to the database definition or create a new database source file.

The **-extract** option does not extract UDF code and Blob filters, because they are not part of the database. It *does* extract the declarations to the database (with DECLARE EXTERNAL FUNCTION and DECLARE FILTER).

The **-extract** option also does not extract system tables, system views, or system triggers.

Because DDL statements do not contain references to object ownership, the extracted file does not show ownership. The output file includes the name of the object and the owner if one is defined. There is no way to assign an object to its original owner.

For a list of the order of extraction of metadata objects, see "Extracting metadata" on page 10-18.For example, the following statement extracts the system catalogs from the database *employee.gdb* to a file called *employee.sql*:

```
isql -extract -output employee.sql employee.gdb;
```

The resulting output script is created with **-commit** following each set of commands, so that tables can be referenced in subsequent definitions. This command extracts all keywords and object names in uppercase when possible (some international metadata has no uppercase).

To extract DDL statements from database *employee.gdb* and store in the file *employee.sql*, enter:

```
isql -a employee.gdb -output employee.sql
```

The following example extracts the DDL statements from the database *dev.ib*:

```
isql -x dev.ib
```

This example combines the **-extract** and **-output** options to extract the DDL statements from the database *dev.ib* into a file called *dev.out*. The output database name must follow the **-output** flag.

```
isql -extract -output dev.out dev.ib
```

isql Commands

At the SQL> prompt, you can enter any of three kinds of commands:

- SQL data definition (DDL) statements, such as CREATE, ALTER, DROP, GRANT, and REVOKE. These statements create, modify, or remove metadata and objects, and control user access (via privileges) to the database. For more information about DDL, see the *Data Definition Guide*.
- SQL data manipulation (DML) statements such as SELECT, INSERT, UPDATE, and DELETE. These four data manipulation operations affect the data in a database. They retrieve, modify, add, or delete data. For more information about DML statements, see the *Language Reference*.
- **isql** commands that fall into three main categories:
  - SHOW commands (to display metadata or other database information)
  - SET commands (to modify the **isql** environment)

- Other commands (for example, commands to read an input file, write to an output file, or end an **isql** session)

Some **isql** commands have many options. See "isql command reference" on page 10-27

SHOW commands

SHOW commands are used to display metadata, including tables, indexes, procedures, and triggers.

SHOW commands list all of the specified objects or give information about a particular object when used with *name*.

SHOW commands operate on a separate transaction from user statements. They run as READ COMMITTED background statements and acknowledge all metadata changes immediately.

SET commands

SET commands enable you to view and change the **isql** environment.

Other isql commands

The remaining **isql** commands perform a variety of useful tasks, including reading a SQL file, executing shell commands, and exiting **isql**. The other **isql** commands are: BLOBDUMP, EDIT, EXIT, HELP, INPUT, OUTPUT, QUIT, SHELL.

Exiting isql

To exit the **isql** utility and roll back all uncommitted work, enter:

```
SQL> QUIT;
```

To exit the **isql** utility and commit all work, enter:

```
SQL> EXIT;
```

Error handling

InterBase handles errors in **isql** and DSQL in the same way. To indicate the causes of an error, **isql** uses the SQLCODE variable and the InterBase status array.

The following table lists values that are returned to SQLCODE:

Table 10.10 SQLCODE and message summary

SQLCODE	Message	Meaning
< 0	SQLERROR	Error occurred; statement did not execute
0	SUCCESS	Successful execution
+1–99	SQLWARNING	System warning or informational message
+100	NOT FOUND	No qualifying rows found, or end of current active set of rows reached

For a detailed discussion of error handling, see the *Embedded SQL Guide*. For a complete listing of SQLCODE and InterBase status array codes, see the *Language Reference*.

isql command reference

This chapter describes the syntax and usage for commands available only in InterBase **isql** (interactive SQL). These commands are also available in SQL scripts. For a description of the standard DSQL commands available in **isql**, see the *Language Reference*.

Command-line **isql** supports the following special commands:

Table 10.11 isql commands

BLOBDUMP	SET BLOBDISPLAY	SHELL	SHOW INDEX
EDIT	SET COUNT	SHOW CHECK	SHOW INDICES
EXIT	SET ECHO	SHOW DATABASE	SHOW PROCEDURES
HELP	SET LIST	SHOW DOMAINS	SHOW ROLES
INPUT	SET NAMES	SHOW EXCEPTIONS	SHOW SYSTEM
OUTPUT	SET PLAN	SHOW FILTERS	SHOW TABLES
QUIT	SET STATS	SHOW FUNCTIONS	SHOW TRIGGERS
SET	SET TERM	SHOW GENERATORS	SHOW VERSION
SET AUTODDL	SET TIME	SHOW GRANT	SHOW VIEWS

BLOBDUMP

Places the contents of a BLOB column in a named file for reading or editing.

Syntax `BLOBDUMP blob_id filename;`

Argument	Description
blob_id	System-assigned hexadecimal identifier, made up of two hexadecimal numbers separated by a colon (:) • First number is the ID of the table containing the BLOB column • Second number is a sequential number identifying a particular instance of Blob data
filename	Name of the file into which to place Blob contents

Description BLOBDUMP stores Blob data identified by *blob_id* in the file specified by *filename*. Because binary files cannot be displayed, BLOBDUMP is useful for viewing or editing binary data. BLOBDUMP is also useful for saving blocks of text (Blob data) to a file.

To determine the blob_id to supply in the BLOBDUMP statement, issue any SELECT statement that selects a column of Blob data. When the table's columns appear, any Blob columns contain hexadecimal Blob IDs. The display of Blob output can be controlled using SET BLOBDISPLAY.

Example Suppose that Blob ID 58:c59 refers to graphical data in JPEG format. To place this Blob data into a graphics file named *picture.jpg*, enter:

```
BLOBDUMP 58:c59 picture.jpg;
```

See also SET BLOBDISPLAY

EDIT

Allows editing and re-execution of **isql** commands.

Syntax `EDIT [filename];`

Argument	Description
filename	Name of the file to edit

Description The EDIT command enables you to edit commands in:

- A source file and then execute the commands upon exiting the editor.
- The current **isql** session, then re-execute them.

On Windows platforms, EDIT calls the text editor specified by the EDITOR environment variable. If this environment variable is not defined, then EDIT uses the Microsoft mep editor.

On UNIX, EDIT calls the text editor specified by either the VISUAL environment variable or EDITOR, in that order. If neither variable is defined, then EDIT uses the **vi** editor.

If given filename as an argument, EDIT places the contents of filename in an edit buffer. If no file name is given, EDIT places the commands in the current isql session in the edit buffer.

After exiting the editor, isql automatically executes the commands in the edit buffer.

Filenames with spaces You can optionally delimit the filename with double or single quotes. This allows you to use filenames with spaces in EDIT statements.

Examples To edit the commands in a file called *start.sql* and execute the commands when done, enter:

```
EDIT START.SQL;
```

In the next example, a user wants to enter `SELECT DISTINCT JOB_CODE, JOB_TITLE FROM JOB;` interactively: Instead, the user mistakenly omits the DISTINCT keyword. Issuing the `EDIT` command opens the statement in an editor and then executes the edited statement when the editor exits.

```
SELECT JOB_CODE, JOB_TITLE FROM JOB;
EDIT;
```

See also INPUT, OUTPUT, SHELL

EXIT

Commits the current transaction, closes the database, and ends the **isql** session.

Syntax `EXIT;`

Description Both EXIT and QUIT close the database and end an **isql** session. EXIT commits any changes made since the last COMMIT or ROLLBACK, whereas QUIT rolls them back.

EXIT is equivalent to the end-of-file character, which differs across systems.

Important EXIT commits changes without prompting for confirmation. Before using EXIT, be sure that no transactions need to be rolled back.

See also QUIT, SET AUTODDL

HELP

Displays a list of ISQL commands and short descriptions.

Syntax `HELP;`

Description HELP lists the built-in **isql** commands, with a brief description of each.

Example To save the HELP screen to a file named *isqlhelp.lst*, enter:

```
OUTPUT isqlhelp.lst;
HELP;
OUTPUT;
```

After issuing the HELP command, use OUTPUT to redirect output back to the screen.

INPUT

Read and execute commands from the named file.

Syntax `INPUT filename;\`

Argument	Description
filename	Name of the file containing SQL statements and SQL commands

Description INPUT reads commands from *filename* and executes them as a block. In this way, INPUT enables execution of commands without prompting. *filename* must contain SQL statements or **isql** commands.

Input files can contain their own INPUT commands. Nesting INPUT commands enables **isql** to process multiple files. When **isql** reaches the end of one file, processing returns to the previous file until all commands are executed.

The INPUT command is intended for noninteractive use. Therefore, the EDIT command does not work in input files.

Using INPUT *filename* from within an **isql** session has the same effect as using **-input *filename*** from the command line.

Unless output is redirected using OUTPUT, any results returned by executing filename appear on the screen.

You can optionally delimit the filename with double or single quotes. This allows you to use filenames with spaces in INPUT statements.

Examples For this example, suppose that file *add.lst* contains the following INSERT statement:

```
INSERT INTO COUNTRY (COUNTRY, CURRENCY)
VALUES ('Mexico', 'Peso');
```

To execute the command stored in *add.lst*, enter:

```
INPUT add.lst;
```

For the next example, suppose that the file, *table.lst*, contains the following SHOW commands:

```
SHOW TABLE COUNTRY;
SHOW TABLE CUSTOMER;
SHOW TABLE DEPARTMENT;
SHOW TABLE EMPLOYEE;
SHOW TABLE EMPLOYEE_PROJECT;
SHOW TABLE JOB;
```

To execute these commands, enter:

```
INPUT table.lst;
```

To record each command and store its results in a file named *table.out*, enter

```
SET ECHO ON;
OUTPUT table.out;
INPUT table.lst;
OUTPUT;
```

See also OUTPUT

OUTPUT

Redirects output to the named file or to standard output.

Syntax `OUTPUT [filename];`

Argument	Description
filename	Name of the file in which to save output; if no file name is given, results appear on the standard output

Description OUTPUT determines where the results of **isql** commands are displayed. By default, results are displayed on standard output (usually a screen). To store results in a file, supply a *filename* argument. To return to the default mode, again displaying results on the standard output, use OUTPUT without specifying a file name.

By default, only data is redirected. Interactive commands are not redirected unless SET ECHO is in effect. If SET ECHO is in effect, **isql** displays each command before it is executed. In this way, **isql** captures both the results and the command that produced them. SET ECHO is useful for displaying the text of a query immediately before the results.

Note Error messages cannot be redirected to an output file.

Using OUTPUT *filename* from within an **isql** session has the same effect as using the option **-output *filename*** from the command line.

You can optionally delimit the filename with double or single quotes. This allows you to use filenames with spaces in OUTPUT statements.

Example The following example stores the results of one SELECT statement in the file, *sales.out*. Normal output processing resumes after the SELECT statement.

```
OUTPUT sales.out;
SELECT * FROM SALES;
OUTPUT;
```

See also INPUT, SET ECHO

QUIT

Rolls back the current transaction, closes the database, and ends the **isql** session.

Syntax

```
QUIT;
```

Description Both EXIT and QUIT close the database and end an **isql** session. QUIT rolls back any changes made since the last COMMIT or ROLLBACK, whereas EXIT commits the changes.

Important QUIT rolls back uncommitted changes without prompting for confirmation. Before using QUIT, be sure that any changes that need to be committed are committed. For example, if SET AUTODDL is off, DDL statements must be committed explicitly.

See also EXIT, SET AUTODDL

SET

Lists the status of the features that control an **isql** session.

Syntax

```
SET;
```

Description **isql** provides several SET commands for specifying how data is displayed or how other commands are processed.

The SET command, by itself, verifies which features are currently set. Some SET commands turn a feature on or off. Other SET commands assign values.

Many **isql** SET commands have corresponding SQL statements that provide similar or identical functionality. In addition, some of the **isql** features controlled by SET commands can also be controlled using **isql** command-line options. SET Statements are used to configure the **isql** environment from a script file. Changes to the session setting from SET statements in a script affect the session only while the script is running. After a script completes, the session settings prior to running the script are restored.

The **isql** SET statements are:

Table 10.12 SET statements

Statement	Description	Default
SET AUTODDL	Toggles the commit feature for DDL statements	ON
SET BLOBDISPLAY *n*	Turns on the display of Blob type *n*; the parameter *n* is required to display Blob types	OFF
SET COUNT	Toggles the count of selected rows on or off	OFF
SET ECHO	Toggles the display of each command on or off	OFF
SET LIST *string*	Displays columns vertically or horizontally	OFF
SET NAMES	Specifies the active character set	OFF
SET PLAN	Specifies whether or not to display the optimizer's query plan	OFF
SET STATS	Toggles the display of performance statistics on or off	OFF
SET TERM *string*	Allows you to change to an alternate terminator character (deprecated in InterBase 7)	;
SET TIME	Toggles display of time in DATE values	ON

By default, all settings are initially OFF except AUTODDL and TIME, and the terminator is a semicolon (;). Each time you start an **isql** session or execute an **isql** script file, settings begin with their default values.

SET statements are used to configure the **isql** environment from a script file. Changes to the session setting from SET statements in a script affect the session only while the script is running. After a script completes, the session settings prior to running the script are restored to their values before the script was run. So you can modify the settings for interactive use, then change them as needed in an **isql** script, and after running the script they automatically return to their previous configuration.

Notes

- You cannot enter **isql** SET statements interactively in the SQL Statement area of IBConsole ISQL. You can perform the same functions with menu items.
- SET GENERATOR and SET TRANSACTION (without a transaction name) are DSQL statements and so you can enter them interactively in IBConsole ISQL or **isql**. These statements are not exclusively **isql** statements, so they are not documented in this chapter. See the *Language Reference* for details.

- SET DATABASE is exclusively an embedded SQL statement. See the *Language Reference* and the *Embedded SQL Guide* for details.

Example To display the **isql** features currently in effect, enter:

```
SET;
   Print statistics:OFF
   Echo commands:OFF
   List format:  OFF
   Row count:    OFF
   Autocommit DDL:OFF
   Access plan:  OFF
   Display BLOB type:1
   Terminator:   ;
   Time:         OFF
```

The output shows that **isql** is set to not echo commands, to display Blob data if they are of subtype 1 (text), to automatically commit DDL statements, and to recognize a semicolon (;) as the statement termination character.

See also SET AUTODDL, SET BLOBDISPLAY, SET COUNT, SET ECHO, SET LIST, SET NAMES, SET PLAN, SET STATS, SET TIME

SET AUTODDL

Specifies whether DDL statements are committed automatically after being executed or committed only after an explicit COMMIT.

Syntax

```
SET AUTODDL [ON | OFF];
```

Argument	Description
ON	Turns on automatic commitment of DDL [default]
OFF	Turns off automatic commitment of DDL

Description SET AUTODDL is used to turn on or off the automatic commitment of data definition language (DDL) statements. By default, DDL statements are automatically committed immediately after they are executed, in a separate transaction. This is the recommended behavior.

If the OFF keyword is specified, auto-commit of DDL is then turned off. In OFF mode, DDL statements can only be committed explicitly through a user's transaction. This mode is useful for database prototyping, because uncommitted changes are easily undone by rolling them back.

SET AUTODDL has a shorthand equivalent, SET AUTO.

Tip The ON and OFF keywords are optional. If they are omitted, SET AUTO switches from one mode to the other. Although you can save typing by omitting the optional keyword, including the keyword is recommended because it avoids potential confusion.

Examples The following example shows part of an **isql** script that turns off AUTODDL, creates a table named TEMP, then rolls back the work.

```
. . .
SET AUTO OFF;
CREATE TABLE TEMP (a INT, b INT);
ROLLBACK;
. . .
```

This script creates TEMP and then rolls back the statement. No table is created. because its creation was rolled back.

The next script uses the default AUTODDL ON. It creates the table TEMP and then performs a rollback:

```
. . .
CREATE TABLE TEMP (a INT, b INT);
ROLLBACK;
. . .
```

Because DDL is automatically committed, the rollback does not affect the creation of TEMP.

See also EXIT, QUIT

SET BLOBDISPLAY

Specifies subtype of Blob data to display.

Syntax `SET BLOBDISPLAY [n | ALL | OFF];`

Argument	Description
n	Integer specifying the Blob subtype to display • Use 0 for Blob data of an unknown subtype • Use 1 for Blob data of a text subtype [default] • Use other integer values for other subtypes
ALL	Displays Blob data of all subtypes
OFF	Turns off display of Blob data of all subtypes

Description SET BLOBDISPLAY has the following uses:

- To display Blob data of a particular subtype, use SET BLOBDISPLAY *n*. By default, **isql** displays Blob data of text subtype (*n* = 1).

- To display Blob data of all subtypes, use SET BLOBDISPLAY ALL.
- To avoid displaying Blob data, use SET BLOBDISPLAY OFF. Omitting the OFF keyword has the same effect. Turn Blob display off to make output easier to read.

In any column containing Blob data, the actual data does not appear in the column. Instead, the column displays a Blob ID that represents the data. If SET BLOBDISPLAY is on, data associated with a Blob ID appears under the row containing the Blob ID. If SET BLOBDISPLAY is off, the Blob ID still appears even though its associated data does not.

SET BLOBDISPLAY has a shorthand equivalent, SET BLOB.

To determine the subtype of a BLOB column, use SHOW TABLE.

Examples The following examples show output from the same SELECT statement. Each example uses a different SET BLOB command to affect how output appears. The first example turns off Blob display.

```
SET BLOB OFF;
SELECT PROJ_NAME, PROJ_DESC FROM PROJECT;
```

With BLOBDISPLAY OFF, the output shows only the Blob ID:

```
PROJ_NAME          PROJ_DESC
=================== =================
Video Database     24:6
DigiPizza          24:8
AutoMap            24:a
MapBrowser port    24:c
Translator upgrade 24:3b
Marketing project 3 24:3d
```

The next example restores the default by setting BLOBDISPLAY to subtype 1 (text).

```
SET BLOB 1;
SELECT PROJ_NAME, PROJ_DESC FROM PROJECT;
```

Now the contents of the Blob appear below each Blob ID:

```
PROJ_NAME     PROJ_DESC
====================================
Video Database 24:6
================================================================
PROJ_DESC:
Design a video data base management system for
controlling on-demand video distribution.
PROJ_NAME     PROJ_DESC
====================================
DigiPizza     24:8
================================================================
PROJ_DESC:
Develop second generation digital pizza maker
```

```
with flash-bake heating element and
digital ingredient measuring system.
. . .
```

See also BLOBDUMP

SET COUNT

Specifies whether to display number of rows retrieved by queries.

Syntax

```
SET COUNT [ON | OFF];
```

Argument	Description
ON	Turns on display of the "rows returned" message
OFF	Turns off display of the "rows returned" message [default]

Description By default, when a SELECT statement retrieves rows from a query, no message appears to say how many rows were retrieved.

Use SET COUNT ON to change the default behavior and display the message. To restore the default behavior, use SET COUNT OFF.

Tip The ON and OFF keywords are optional. If they are omitted, SET COUNT switches from one mode to the other. Although you can save typing by omitting the optional keyword, including the keyword is recommended because it avoids potential confusion.

Example The following example sets COUNT ON to display the number of rows returned by all following queries:

```
SET COUNT ON;
SELECT * FROM COUNTRY
   WHERE CURRENCY LIKE '%FRANC%';
```

The output displayed would then be:

```
   COUNTRY       CURRENCY
   ============== ==========
   SWITZERLAND   SFRANC
   FRANCE        FFRANC
   BELGIUM       BFRANC
3 rows returned
```

SET ECHO

Specifies whether commands are displayed to the **isql** Output area before being executed.

Syntax

```
SET ECHO [ON | OFF];
```

Argument	Description
ON	Turns on command echoing [default]
OFF	Turns off command echoing

Description

By default, commands in script files are displayed (echoed) in the **isql** Output area, before being executed. Use SET ECHO OFF to change the default behavior and suppress echoing of commands. This can be useful when sending the output of a script to a file, if you want only the results of the script and not the statements themselves in the output file.

Command echoing is useful if you want to see the commands as well as the results in the **isql** Output area.

Tip

The ON and OFF keywords are optional. If they are omitted, SET ECHO switches from one mode to the other. Although you can save typing by omitting the optional keyword, including the keyword is recommended because it avoids potential confusion.

Example

Suppose you execute the following script from IBConsole ISQL:

```
. . .
SET ECHO OFF;
SELECT * FROM COUNTRY;
SET ECHO ON;
SELECT * FROM COUNTRY;
EXIT;
```

The output (in a file or the **isql** Output area) looks like this:

```
. . .
SET ECHO OFF;
COUNTRY    CURRENCY
====================
USA        Dollar
England    Pound
. . .
SELECT * FROM COUNTRY;
COUNTRY    CURRENCY
====================
USA        Dollar
England    Pound
. . .
```

The first SELECT statement is not displayed, because ECHO is OFF. Notice also that the SET ECHO ON statement itself is not displayed, because when it is executed, ECHO is still OFF. After it is executed, however, the second SELECT statement is displayed.

See also INPUT, OUTPUT

SET LIST

Specifies whether output appears in tabular format or in list format.

Syntax

```
SET LIST [ON | OFF];
```

Argument	Description
ON	Turns on list format for display of output
OFF	Turns off list format for display of output [default]

Description By default, when a SELECT statement retrieves rows from a query, the output appears in a tabular format, with data organized in rows and columns.

Use SET LIST ON to change the default behavior and display output in a list format. In list format, data appears one value per line, with column headings appearing as labels. List format is useful when columnar output is too wide to fit nicely on the screen.

Tip The ON and OFF keywords are optional. If they are omitted, SET LIST switches from one mode to the other. Although you can save typing by omitting the optional keyword, including the keyword is recommended because it avoids potential confusion.

Example Suppose you execute the following statement in a script file:

```
SELECT JOB_CODE, JOB_GRADE, JOB_COUNTRY, JOB_TITLE FROM JOB
   WHERE JOB_COUNTRY = 'Italy';
```

The output is:

```
JOB_CODE   JOB_GRADE   JOB_COUNTRY   JOB_TITLE
========   =========   ===========   ====================
SRep       4           Italy         Sales Representative
```

Now suppose you precede the SELECT with SET LIST ON:

```
SET LIST ON;
SELECT JOB_CODE, JOB_GRADE, JOB_COUNTRY, JOB_TITLE FROM JOB
   WHERE JOB_COUNTRY = 'Italy';
```

The output is:

```
   JOB_CODE      SRep
   JOB_GRADE     4
```

```
JOB_COUNTRY   Italy
JOB_TITLE     Sales Representative
```

SET NAMES

Specifies the active character set to use in database transactions.

Syntax

```
SET NAMES [charset];
```

Argument	Description
charset	Name of the active character set; default is NONE

Description

SET NAMES specifies the character set to use for subsequent database connections in **isql**. It enables you to override the default character set for a database. To return to using the default character set, use SET NAMES with no argument.

Use SET NAMES before connecting to the database whose character set you want to specify. For a complete list of character sets recognized by InterBase, see the *Language Reference*.

Choice of character set limits possible collation orders to a subset of all available collation orders. Given a specific character set, a specific collation order can be specified when data is selected, inserted, or updated in a column.

Example

The following statement at the beginning of a script file indicates to set the active character set to ISO8859_1 for the subsequent database connection:

```
SET NAMES ISO8859_1;
CONNECT 'jupiter:/usr/interbase/examples/employee.gdb';
. . .
```

SET PLAN

Specifies whether to display the optimizer's query plan.

Syntax

```
SET PLAN [ON | OFF];
```

Argument	Description
ON	Turns on display of the optimizer's query plan
OFF	Turns off display of the optimizer's query plan [default]

Description

By default, when a SELECT statement retrieves rows from a query, **isql** does not display the query plan used to retrieve the data.

Use SET PLAN ON to change the default behavior and display the query optimizer plan. To restore the default behavior, use SET PLAN OFF.

To change the query optimizer plan, use the PLAN clause in the SELECT statement.

Tip The ON and OFF keywords are optional. If they are omitted, SET PLAN switches from one mode to the other. Although you can save typing by omitting the optional keyword, including the keyword is recommended because it avoids potential confusion.

Example The following example shows part of a script that sets PLAN ON:

```
SET PLAN ON;
SELECT JOB_COUNTRY, MIN_SALARY FROM JOB
   WHERE MIN_SALARY > 50000
      AND JOB_COUNTRY = 'France';
```

The output then includes the query optimizer plan used to retrieve the data as well as the results of the query:

```
PLAN (JOB INDEX (RDB$FOREIGN3,MINSALX,MAXSALX))
JOB_COUNTRY   MIN_SALARY
=============== ======================
France        118200.00
```

SET STATS

Specifies whether to display performance statistics after the results of a query.

Syntax `SET STATS [ON | OFF];`

Argument	Description
ON	Turns on display of performance statistics
OFF	Turns off display of performance statistics [default]

Description By default, when a SELECT statement retrieves rows from a query, **isql** does not display performance statistics after the results. Use SET STATS ON to change the default behavior and display performance statistics. To restore the default behavior, use SET STATS OFF. Performance statistics include:

- Current memory available, in bytes
- Change in available memory, in bytes
- Maximum memory available, in bytes
- Elapsed time for the operation
- CPU time for the operation
- Number of cache buffers used
- Number of reads requested
- Number of writes requested

- Number of fetches made

Performance statistics can help determine if changes are needed in system resources, database resources, or query optimization.

Tip The ON and OFF keywords are optional. If they are omitted, SET STATS switches from one mode to the other. Although you can save typing by omitting the optional keyword, including the keyword is recommended because it avoids potential confusion.

Do not confuse SET STATS with the SQL statement SET STATISTICS, which recalculates the selectivity of an index.

Example The following part of a script file turns on display of statistics and then performs a query:

```
SET STATS ON;
SELECT JOB_COUNTRY, MIN_SALARY FROM JOB
   WHERE MIN_SALARY > 50000
      AND JOB_COUNTRY = 'France';
```

The output displays the results of the SELECT statement and the performance statistics for the operation:

```
   JOB_COUNTRY   MIN_SALARY
   =============== ======================
   France        118200.00

Current memory = 407552
Delta memory = 0
Max memory = 412672
Elapsed time= 0.49 sec
Cpu = 0.06 sec
Buffers = 75
Reads = 3
Writes = 2
Fetches = 441
```

See also SHOW DATABASE

SET TERM

Specifies which character or characters signal the end of a command.

Syntax `SET TERM` *string*;

Argument	Description
string	Specifies a character or characters to use in terminating a statement; default is semicolon (;)

Description By default, when a line ends with a semicolon, **isql** interprets it as the end of a command. Use SET TERM to change the default behavior and define a new termination character.

SET TERM is typically used with CREATE PROCEDURE or CREATE TRIGGER. Procedures and triggers are defined using a special "procedure and trigger language" in which statements end with a semicolon. If **isql** were to interpret semicolons as statement terminators, then procedures and triggers would execute during their creation, rather than when they are called.

A script file containing CREATE PROCEDURE or CREATE TRIGGER definitions should include one SET TERM command before the definitions and a corresponding SET TERM after the definitions. The beginning SET TERM defines a new termination character; the ending SET TERM restores the semicolon (;) as the default.

Note You do not need to change the terminator before entering an interactive CREATE PROCEDURE or CREATE TRIGGER statement in the IBConsole ISQL SQL statement area. The contents of the SQL statement area is always treated as one DSQL statement, even if it contains semicolons. Use of SET TERM is necessary only in command-line **isql** and when running SQL script files from command-line **isql** or IBConsole ISQL.

Example The following example shows a text file that uses SET TERM in creating a procedure. The first SET TERM defines "##" as the termination characters. The matching SET TERM restores ";" as the termination character.

```
SET TERM ## ;
CREATE PROCEDURE ADD_EMP_PROJ (EMP_NO SMALLINT, PROJ_ID CHAR(5))
AS
BEGIN
   BEGIN
      INSERT INTO EMPLOYEE_PROJECT (EMP_NO, PROJ_ID)
         VALUES (:emp_no, :proj_id);
   WHEN SQLCODE -530 DO
   EXCEPTION UNKNOWN_EMP_ID;
   END
   RETURN;
END ##
SET TERM ; ##
```

SET TIME

Specifies whether to display the time portion of a DATE value.

Syntax `SET TIME [ON | OFF];`

Argument	Description
ON	Turns on display of time in DATE value
OFF	Turns off display of time in DATE value [default]

Description The InterBase Date datatype includes a date portion (including day, month, and year) and a time portion (including hours, minutes, and seconds).

By default, **isql** displays only the date portion of Date values. SET TIME ON turns on the display of time values. SET TIME OFF turns off the display of time values.

Tip The ON and OFF keywords are optional. If they are omitted, the command toggles time display from ON to OFF or OFF to ON.

Example The following example shows the default display of a DATE datatype, which is to display day, month, and year:

```
SELECT HIRE_DATE FROM EMPLOYEE WHERE EMP_NO = 145;
HIRE_DATE
==========
2-MAY-1994
```

This example shows the effects of SET TIME ON, which causes the hours, minutes and seconds to be displayed as well:

```
SET TIME ON;
SELECT HIRE_DATE FROM EMPLOYEE WHERE EMP_NO = 145;
HIRE_DATE
==========
2-MAY-1994 12:25:00
```

SHELL

Allows execution of an operating system command or temporary access to an operating system shell.

Syntax `SHELL [<os_command>];`

Argument	Description
os_command	An operating system command; if no command is specified, **isql** provides interactive access to the operating system

Description The SHELL command provides temporary access to operating system commands in an **isql** session. Use SHELL to execute an operating-system command without ending the current **isql** session.

If *os_command* is specified, the operating system executes the command and then returns to **isql** when complete.

If no command is specified, an operating system shell prompt appears, enabling you to execute a sequence of commands. To return to **isql**, type `exit`. For example, SHELL can be used to edit an input file and run it at a later time. By contrast, if an input file is edited using the EDIT command, the input file is executed as soon as the editing session ends.

Using SHELL does not commit transactions before it calls the shell.

This **isql** statement has no equivalent function in IBConsole ISQL.

Example The following example uses SHELL to display the contents of the current directory:

```
SHELL DIR;
```

See also EDIT

SHOW CHECK

Displays all CHECK constraints defined for a specified table.

Syntax `SHOW CHECK table;`

Argument	Description
table	Name of an existing table in the current database

Description SHOW CHECK displays CHECK constraints for a named table in the current database. Only user-defined metadata is displayed. To see a list of existing tables, use SHOW TABLE.

Example The following example shows CHECK constraints defined for the JOB table. The SHOW TABLES command is used first to display a list of available tables.

```
SHOW TABLES;
   COUNTRY       CUSTOMER
   DEPARTMENT    EMPLOYEE
   EMPLOYEE_PROJECTJOB
   PHONE_LIST    PROJECT
   PROJ_DEPT_BUDGETSALARY_HISTORY
   SALES

SHOW CHECK JOB;
   CHECK (min_salary < max_salary)
```

See also SHOW TABLES

SHOW DATABASE

Displays information about the current database.

Syntax `SHOW [DATABASE | DB];`

Description SHOW DATABASE displays the current database's file name, page size and allocation, and sweep interval.

The output of SHOW DATABASE is used to verify data definition or to administer the database. For example, use the backup and restore utilities to change page size or reallocate pages among multiple files, and use the database maintenance utility to change the sweep interval.

SHOW DATABASE has a shorthand equivalent, SHOW DB.

Example The following example connects to a database and displays information about it:

```
CONNECT 'employee.gdb';
   Database: employee.gdb

SHOW DB;
   Database: employee.gdb
        Owner: SYSDBA
   PAGE_SIZE 4096
   Number of DB pages allocated = 422
   Sweep interval = 20000
```

SHOW DOMAINS

Lists all domains or displays information about a specified domain.

Syntax SHOW {DOMAINS | DOMAIN *name*};

Argument	Description
name	Name of an existing domain in the current database

Options To see a list of existing domains, use SHOW DOMAINS without specifying a domain name. SHOW DOMAIN name displays information about the named domain in the current database. Output includes a domain's datatype, default value, and any CHECK constraints defined. Only user-defined metadata is displayed.

Example The following example lists all domains and then shows the definition of the domain, SALARY:

```
SHOW DOMAINS;
   FIRSTNAME     LASTNAME
   PHONENUMBER   COUNTRYNAME
   ADDRESSLINE   EMPNO
   DEPTNO        PROJNO
   CUSTNO        JOBCODE
   JOBGRADE      SALARY
   BUDGET        PRODTYPE
   PONUMBER
```

```
SHOW DOMAIN SALARY;
   SALARY        NUMERIC(15, 2) Nullable
                 DEFAULT 0
                 CHECK (VALUE > 0)
```

SHOW EXCEPTIONS

Lists all exceptions or displays the text of a specified exception.

Syntax `SHOW {EXCEPTIONS | EXCEPTION` *name*`};`

Argument	Description
name	Name of an existing exception in the current database

Description SHOW EXCEPTIONS displays an alphabetical list of exceptions. SHOW EXCEPTION name displays the text of the named exception.

Examples To list all exceptions defined for the current database, enter:

```
SHOW EXCEPTIONS;
Exception Name Used by, Type
================= ======================================
UNKNOWN_EMP_ID ADD_EMP_PROJ, Stored procedure
   Invalid employee number or project ID.
. . .
```

To list the message for a specific exception and the procedures or triggers that use it, enter the exception name:

```
SHOW EXCEPTION CUSTOMER_CHECK;
Exception Name                  Used by, Type
==========================  ======================================
CUSTOMER_CHECK              SHIP_ORDER, Stored procedure
 Overdue balance -- can't ship.
```

SHOW FILTERS

Lists all Blob filters or displays information about a specified filter.

Syntax `SHOW {FILTERS | FILTER` *name*`};`

Argument	Description
name	Name of an existing Blob filter in the current database

Options To see a list of existing filters, use SHOW FILTERS. SHOW FILTER name displays information about the named filter in the current database. Output includes information previously defined by the DECLARE FILTER statement, the input subtype, output subtype, module (or library) name, and entry point name.

Example The following example lists all filters and then shows the definition of the filter, DESC_FILTER:

```
SHOW FILTERS;
   DESC_FILTER

SHOW FILTER DESC_FILTER;
   BLOB Filter: DESC_FILTER
   Input subtype: 1 Output subtype -4
   Filter library is: desc_filter
   Entry point is: FILTERLIB
```

SHOW FUNCTIONS

Lists all user-defined functions (UDFs) defined in the database or displays information about a specified UDF.

Syntax `SHOW {FUNCTIONS | FUNCTION` *name*`};`

Argument	Description
name	Name of an existing UDF in the current database

Options To see a list of existing functions defined in the database, use SHOW FUNCTIONS. SHOW FUNCTION name displays information about the named function in the current database. Output includes information previously defined by the DECLARE EXTERNAL FUNCTION statement: the name of the function and function library, the name of the entry point, and the datatypes of return values and input arguments.

Example The following example lists all UDFs and then shows the definition of the MAXNUM() function:

```
SHOW FUNCTIONS;
   ABS      MAXNUM
   TIME     UPPER_NON_C
   UPPER

SHOW FUNCTION maxnum;
   Function MAXNUM:
   Function library is /usr/interbase/lib/gdsfunc.so
   Entry point is FN_MAX
   Returns BY VALUE DOUBLE PRECISION
   Argument 1: DOUBLE PRECISION
   Argument 2: DOUBLE PRECISION
```

SHOW GENERATORS

Lists all generators or displays information about a specified generator.

Syntax `SHOW {GENERATORS | GENERATOR name};`

Argument	Description
name	Name of an existing generator in the current database

Description To see a list of existing generators, use SHOW GENERATORS. SHOW GENERATOR name displays information about the named generator in the current database. Output includes the name of the generator and its next value.

SHOW GENERATOR has a shorthand equivalent, SHOW GEN.

Example The following example lists all generators and then shows information about EMP_NO_GEN:

```
SHOW GENERATORS;
   Generator EMP_NO_GEN, Next value: 146
   Generator CUST_NO_GEN, Next value: 1016

SHOW GENERATOR EMP_NO_GEN;
   Generator EMP_NO_GEN, Next value: 146
```

SHOW GRANT

Displays privileges for a database object.

Syntax `SHOW GRANT object;`

Argument	Description
object	Name of an existing table, view, or procedure in the current database

Description SHOW GRANT displays the privileges defined for a specified table, view, or procedure. Allowed privileges are DELETE, EXECUTE, INSERT, SELECT, UPDATE, or ALL. To change privileges, use the SQL statements GRANT or REVOKE.

Before using SHOW GRANT, you might want to list the available database objects. Use SHOW PROCEDURES to list existing procedures; use SHOW TABLES to list existing tables; use SHOW VIEWS to list existing views.

Example To display GRANT privileges on the JOB table, enter:

```
SHOW GRANT JOB;
   GRANT SELECT ON JOB TO ALL
   GRANT DELETE, INSERT, SELECT, UPDATE ON JOB TO MANAGER
```

SHOW GRANT can also show role membership:

```
SHOW GRANT DOITALL;
   GRANT DOITALL TO SOCKS
```

See also SHOW PROCEDURES, SHOW TABLES, SHOW VIEWS

SHOW INDEX

Displays index information for a specified index, for a specified table, or for all tables in the current database.

Syntax `SHOW {INDICES | INDEX {index | table} };`

Argument	Description
index	Name of an existing index in the current database
table	Name of an existing table in the current database

Description SHOW INDEX displays the index name, the index type (for example, UNIQUE or DESC), and the columns on which an index is defined.

If the index argument is specified, SHOW INDEX displays information only for that index. If table is specified, SHOW INDEX displays information for all indexes in the named table; to display existing tables, use SHOW TABLES. If no argument is specified, SHOW INDEX displays information for all indexes in the current database.

SHOW INDEX has a shorthand equivalent, SHOW IND. SHOW INDICES is also a synonym for SHOW INDEX. SHOW INDEXES is not supported.

Examples To display indexes for database *employee.gdb*, enter:

```
SHOW INDEX;
   RDB$PRIMARY1 UNIQUE INDEX ON COUNTRY(COUNTRY)
   CUSTNAMEX INDEX ON CUSTOMER(CUSTOMER)
   CUSTREGION INDEX ON CUSTOMER(COUNTRY, CITY)
   RDB$FOREIGN23 INDEX ON CUSTOMER(COUNTRY)
. . .
```

To display index information for the SALES table, enter:

```
SHOW IND SALES;
   NEEDX INDEX ON SALES(DATE_NEEDED)
   QTYX DESCENDING INDEX ON SALES(ITEM_TYPE, QTY_ORDERED)
   RDB$FOREIGN25 INDEX ON SALES(CUST_NO)
   RDB$FOREIGN26 INDEX ON SALES(SALES_REP)
   RDB$PRIMARY24 UNIQUE INDEX ON SALES(PO_NUMBER)
   SALESTATX INDEX ON SALES(ORDER_STATUS, PAID)
```

See also SHOW TABLES

SHOW PROCEDURES

Lists all procedures or displays the text of a specified procedure.

Syntax

```
SHOW {PROCEDURES | PROCEDURE name};
```

Argument	Description
name	Name of an existing procedure in the current database

Description

SHOW PROCEDURES displays an alphabetical list of procedures, along with the database objects they depend on. Deleting a database object that has a dependent procedure is not allowed. To avoid an **isql** error, delete the procedure (using DROP PROCEDURE) before deleting the database object.

SHOW PROCEDURE name displays the text and parameters of the named procedure.

SHOW PROCEDURE has a shorthand equivalent, SHOW PROC.

Examples

To list all procedures defined for the current database, enter:

```
SHOW PROCEDURES;

   Procedure Name Dependency Type
   ================ ================== =======
   ADD_EMP_PROJ  EMPLOYEE_PROJECTTable
                 UNKNOWN_EMP_IDException
   DELETE_EMPLOYEE DEPARTMENTTable
                 EMPLOYEETable
                 EMPLOYEE_PROJECTTable
                 PROJECTTable
                 REASSIGN_SALESException
                 SALARY_HISTORYTable
                 SALES Table
   DEPT_BUDGET   DEPARTMENTTable
                 DEPT_BUDGETProcedure
. . .
```

To display the text of the procedure, ADD_EMP_PROJ, enter:

```
SHOW PROC ADD_EMP_PROJ;

   Procedure text:
   ========================================================================

   BEGIN
      BEGIN
      INSERT INTO EMPLOYEE_PROJECT (EMP_NO, PROJ_ID) VALUES (:emp_no,
         :proj_id);
      WHEN SQLCODE -530 DO
      EXCEPTION UNKNOWN_EMP_ID;
```

```
    END
    RETURN;
END
============================================================
Parameters:
EMP_NO INPUT SMALLINT
PROJ_ID INPUT CHAR(5)
```

SHOW ROLES

Displays the names of SQL roles for the current database.

Syntax `SHOW {ROLES | ROLE}`

Description SHOW ROLES displays the names of all roles defined for the current database. To show user membership in roles, use SHOW GRANT *rolename*.

Example

```
SHOW ROLES;
   DOITALL    DONOTHING
   DOONETHING DOSOMETHING
```

See also SHOW GRANT

SHOW SYSTEM

Displays the names of system tables and system views for the current database.

Syntax `SHOW SYSTEM [TABLES];`

Description SHOW SYSTEM lists system tables and system views in the current database. SHOW SYSTEM accepts an optional keyword, TABLES, which does not affect the behavior of the command.

SHOW SYSTEM has a shorthand equivalent, SHOW SYS.

Example To list system tables and system views for the current database, enter:

```
SHOW SYS;
   RDB$CHARACTER_SETS RDB$CHECK_CONSTRAINTS
   RDB$COLLATIONS     RDB$DATABASE
   RDB$DEPENDENCIES   RDB$EXCEPTIONS
   RDB$FIELDS         RDB$FIELD_DIMENSIONS
   RDB$FILES          RDB$FILTERS
   RDB$FORMATS        RDB$FUNCTIONS
   RDB$FUNCTION_ARGUMENTS RDB$GENERATORS
   RDB$INDEX_SEGMENTS RDB$INDICES
   RDB$LOG_FILES      RDB$PAGES
   RDB$PROCEDURES     RDB$PROCEDURE_PARAMETERS
   RDB$REF_CONSTRAINTS RDB$RELATIONS
   RDB$RELATION_CONSTRAINTS RDB$RELATION_FIELDS
   RDB$ROLES          RDB$SECURITY_CLASSES
```

```
   RDB$TRANSACTIONS   RDB$TRIGGERS
   RDB$TRIGGER_MESSAGESRDB$TYPES
   RDB$USER_PRIVILEGES RDB$VIEW_RELATIONS
```

See also For more information about system tables, see the *Language Reference*.

SHOW TABLES

Lists all tables or views, or displays information about a specified table or view.

Syntax

```
SHOW {TABLES | TABLE name};
```

Argument	Description
name	Name of an existing table or view in the current database

Description SHOW TABLES displays an alphabetical list of tables and views in the current database. To determine which listed objects are views rather than tables, use SHOW VIEWS.

SHOW TABLE name displays information about the named object. If the object is a table, command output lists column names and definitions, PRIMARY KEY, FOREIGN KEY, and CHECK constraints, and triggers. If the object is a view, command output lists column names and definitions, as well as the SELECT statement that the view is based on.

Examples To list all tables or views defined for the current database, enter:

```
SHOW TABLES;
   COUNTRY       CUSTOMER
   DEPARTMENT    EMPLOYEE
   EMPLOYEE_PROJECT JOB
   PHONE_LIST    PROJECT
   PROJ_DEPT_BUDGET SALARY_HISTORY
   SALES
```

To show the definition for the COUNTRY table, enter:

```
SHOW TABLE COUNTRY;
   COUNTRY (COUNTRYNAME) VARCHAR(15) NOT NULL
   CURRENCY VARCHAR(10) NOT NULL
   PRIMARY KEY (COUNTRY)
```

See also SHOW VIEWS

SHOW TRIGGERS

Lists all triggers or displays information about a specified trigger.

Syntax

```
SHOW {TRIGGERS | TRIGGER name};
```

Argument	Description
name	Name of an existing trigger in the current database

Description SHOW TRIGGERS displays all triggers defined in the database, along with the table they depend on. SHOW TRIGGER name displays the name, sequence, type, activation status, and definition of the named trigger.

SHOW TRIGGER has a shorthand equivalent, SHOW TRIG.

Deleting a table that has a dependent trigger is not allowed. To avoid an **isql** error, delete the trigger (using DROP TRIGGER) before deleting the table.

Examples To list all triggers defined for the current database, enter:

```
SHOW TRIGGERS;
   Table name    Trigger name
   ===========   ============
   EMPLOYEE      SET_EMP_NO
   EMPLOYEE      SAVE_SALARY_CHANGE
   CUSTOMER      SET_CUST_NO
   SALES         POST_NEW_ORDER
```

To display information about the SET_CUST_NO trigger, enter:

```
SHOW TRIG SET_CUST_NO;

   Triggers:
   SET_CUST_NO, Sequence: 0, Type: BEFORE INSERT, Active
   AS
   BEGIN
      new.cust_no = gen_id(cust_no_gen, 1);
   END
```

SHOW VERSION

Displays information about software versions.

Syntax `SHOW VERSION;`

Description SHOW VERSION displays the software version of **isql**, the InterBase engine, and the on-disk structure (ODS) of the database to which the session is attached.

Certain tasks might not work as expected if performed on databases that were created using older versions of InterBase. To check the versions of software that are running, use SHOW VERSION.

SHOW VERSION has a shorthand equivalent, SHOW VER.

Example To display software versions, enter:

```
SQL> SHOW VERSION;
```

```
ISQL Version: WI-V7.0.0.168
InterBase/x86/Windows NT (access method), version "WI-V7.0.0.168"
on disk structure version 11.0
```

See also SHOW DATABASE

SHOW VIEWS

Lists all views or displays information about a specified view.

Syntax `SHOW {VIEWS | VIEW name};`

Argument	Description
name	Name of an existing view in the current database

Description SHOW VIEWS displays an alphabetical list of all views in the current database. SHOW VIEW name displays information about the named view.

Example To list all views defined for the current database, enter:

```
SHOW VIEWS;
   PHONE_LIST
```

See also SHOW TABLES

Using SQL scripts

The basic steps for using script files are:

1. Create the script file using a text editor.
2. Run the file with **isql** or IBConsole.
3. View output and confirm database changes.

Creating an isql script

You can use any text editor to create a SQL script file, as long as the final file format is plain text (ASCII).

Every SQL script file must begin with either a CREATE DATABASE statement or a CONNECT statement (including username and password) that specifies the database on which the script file is to operate. The CONNECT or CREATE statement must contain a complete database file name and directory path.

Note You cannot set dialect in a CREATE DATABASE statement. To create a dialect 3 database, specify **isql** option -s 3.

A SQL script can contain any of the following elements:

- SQL statements, as described in the *Language Reference*
- **isql** SET commands as described in this chapter
- Comments.

Each SQL statement in a script must end with a terminator.

Note The SQL statement silently fails if significant text follows the terminator character on the same line. Whitespace and comments can safely follow the terminator, but other statements cannot.

Each SQL script file should end with either EXIT to commit database changes made since the last COMMIT, or QUIT to roll back changes made by the script. If neither is specified, then database changes are committed by default.

For the full syntax of CONNECT and CREATE DATABASE, see the *Language Reference*.

Running a SQL script

The following steps execute all the SQL statements in the specified script file. The contents of the script are not displayed in the SQL Input Area.

To run a script file containing SQL statements using IBConsole:

1. If you are not already in the SQL window, click the Launch SQL toolbar button or choose **Tools | Interactive SQL**.
2. If you are not running the SQL script on the database to which you are currently connected, then check that the file begins with a valid, uncommented, CONNECT or CREATE DATABASE statement.
3. Choose **Query | Load Script**.
4. Enter or locate the desired script filename in the Open dialog, and click Open to load the script into the SQL input area.
5. Click the Execute toolbar button, or choose **Query | Execute**.

If IBConsole encounters an error, an information dialog appears indicating the error. Once IBConsole finishes executing the script, the script results are displayed in the SQL output window.

After a script executes, all **isql** session settings prior to executing the script are restored as well as the previous database connection, if any. In other words, any **isql** SET commands in the script affect only the **isql** session while the script is running.

Committing work in a SQL script

Changes to the database from data definition (DDL) statements—for example, CREATE and ALTER statements—are automatically committed by default. This means that other users of the database see changes as soon as each DDL statement

is executed. To turn off automatic commit of DDL in a script, use SET AUTODDL OFF, or set it in the Query Options dialog. See "Using InterBase Manager to start and stop InterBase" on page 3-5 for more information.

Note When creating tables and other database objects with AUTODDL OFF, it is good practice to put a COMMIT statement in the SQL script after each CREATE statement or group of related statements. This ensures that other users of the database see the objects immediately.

Changes made to the database by data manipulation (DML) statements—for example INSERT and UPDATE—are not permanent until they are committed. Commit changes in a script with COMMIT. To undo all database changes since the last COMMIT, use ROLLBACK. For the full syntax of COMMIT and ROLLBACK, see the *Language Reference* book.

Adding comments in an isql script

isql scripts are commented exactly like C programs:

```
/* comment */
```

A comment can occur on the same line as a SQL statement or **isql** command and can be of any length, as long as it is preceded by "/*" and followed by "*/".

Chapter

11

Database and Server Performance

This chapter describes techniques for designing and operating an InterBase client/server system for best speed and efficiency.

The guidelines in this chapter are organized into the following categories:

- Hardware configuration
- Operating system configuration
- Network configuration
- Database properties
- Database design principles
- Database tuning tasks
- Application design techniques
- Application development tools

Introduction

One of the most important requirements for a database as part of your application is to store and retrieve data as quickly as possible. Like any software development technique, there is always more than one method to implement a given specified software solution, and it takes knowledge and experience to choose the design that results in the most efficient operation and the highest performance.

Each project offers unique challenges and requires specific solutions. The suggestions in this chapter augment your own software engineering discipline, which should include careful analysis, testing, and experimentation to implement the best design for your specific project.

Hardware configuration

This section gives guidelines for platform hardware sizing. The suggestions focus on requirements for a server platform.

Choosing a processor speed

The performance of database systems tends by nature to be bound by I/O bandwidth or network bandwidth. An application often waits for I/O or network operations, instead of being computationally intensive. A fast CPU clock speed gives definite performance advantage, but a 10% increase in CPU clock speed is less important for server performance than some other hardware factors, such as RAM configuration, I/O system, or network hardware.

CPU clock speed is often more important on client platforms, because applications that use data might perform CPU-intensive computational analysis on data, or might render sophisticated visualization of data in a computationally costly manner.

It's not appropriate for this document to recommend a specific CPU clock speed for your server, because it is likely that such a recommendation would be obsolete as you read it. You should evaluate the benefit of spending more money on a faster CPU, because the price/performance curve becomes steep for the latest CPU hardware.

Sizing memory

It is important to equip your server with a sufficient amount of physical memory to ensure good performance.

While InterBase can function in a low-profile hardware configuration, with as little as 32MB of RAM on most operating systems, it is recommended to have at least 64MB of RAM on a server system. Database servers that experience a high load can benefit from more RAM.

The base RAM requirement of the **ibserver** executable and for each connected user is low: approximately 1500KB, plus 28KB for each client connection. **ibserver** caches metadata and data for each database to which it connects. User operations such as sorting temporarily consume additional memory. A heavily loaded server with dozens of clients performing concurrent queries requires up to 256MB of RAM.

On Windows NT, you can use the Task Manager, Performance Monitor, and other tools to monitor the resource use of **ibserver**. UNIX and Linux servers have similar resource consumption reporting tools. Add RAM to a system that shows too many page faults.

Using high-performance I/O subsystems

A multiuser database server's hard drives are no place to be thrifty, especially in today's market of inexpensive storage. Configuring a relatively high-end I/O system is a cost-effective way to increase performance.

Slow disk subsystems are often the weak link in an otherwise high-performance server machine. The top-rated CPU and maximum memory helps. But if a cheap disk I/O interface limits the data transfer rate, then the money spent on the expensive components is wasted.

It's not appropriate for this document to recommend a particular configuration. The technology changes so quickly that any recommendation here would be outdated. When you specify the machine for a server platform, research the best hardware solution available.

Read the following guidelines for principles:

- Advanced SCSI technology offers superior I/O throughput. The following graph illustrates the relative maximum throughput of different disk interfaces.

Figure 11.1 Comparing external transfer rate of disk I/O interfaces

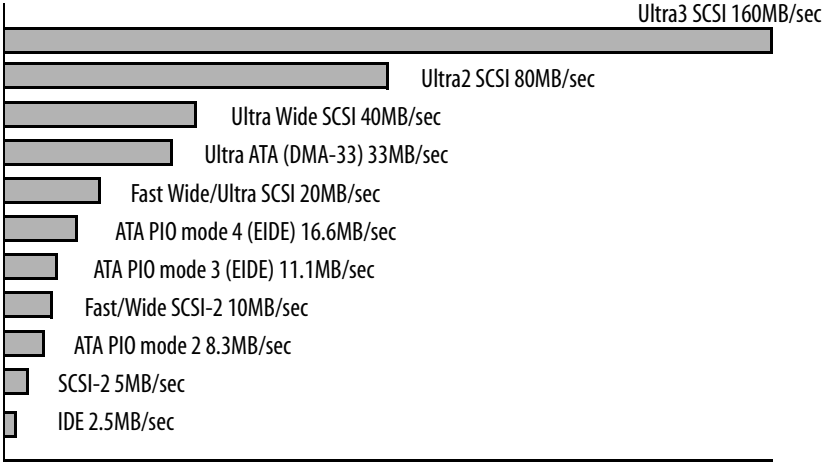

- The external interface capacity usually exceeds the internal or sustained transfer rate of any individual device. Only systems that use multiple disk devices make full use of a high-capacity I/O interface.
- Bus-mastering I/O controllers use less CPU resources. This is particularly important on I/O-intensive server machines. SCSI is generally bus-mastering, and newer PCI EIDE interfaces are bus-mastering. IDE is not.
- Use a disk controller with built in cache memory. The controller cache reduces the need for the operating system to use system RAM for disk cache.

- Don't assume all disks of a given size perform equally; research performance ratings made by independent testing labs.

Distributing I/O

Disk device I/O is orders of magnitude slower than physical memory accesses or CPU cycles. There is a delay while the disk device seeks the data requested. While an application is waiting for data it has requested from a disk device, it is advantageous for the application to spend the time executing other tasks. One appropriate way to do this is to spread multiple data requests over multiple devices. While one disk is preparing to return data, the application requests another disk to start seeking another set of data. This is called *distributed I/O* or *parallel I/O*.

This section describes ways you can persuade InterBase to distribute I/O over multiple disk devices.

Using RAID

You can achieve up to a ten times performance improvement by using RAID.

RAID (redundant array of inexpensive disks) is a hardware design that is intended to give benefits to performance and reliability by storing data on multiple physical disk devices. It is transparent for software applications to use RAID, because it is implemented in the operating system or at the hardware level. InterBase uses operating system I/O interfaces, so InterBase supports RAID as would any other application software.

Disk striping (included in RAID levels 0, 3, or 5) provides performance benefits by distributing I/O across multiple disks.

Hardware RAID is faster than software RAID or software disk mirroring. RAID implemented with software provides only protection from hard disk failure; it is actually slower than operating without RAID.

Using multiple disks for database files

Similarly to RAID, you can distribute files of a multifile InterBase database among multiple physical disk drives.

For example, if you have a server with four physical disks, *C:*, *D:*, *E:*, and *F:*, and a 10GB database, you can create your database to take advantage of parallel I/O with the following database creation statement:

```
CREATE DATABASE 'C:\data\bigdata1.ib' PAGE_SIZE 4096
   FILE 'D:\data\bigdata2.ib' STARTING AT PAGE 1000000
   FILE 'E:\data\bigdata3.ib' STARTING AT PAGE 2000000
   FILE 'F:\data\bigdata4.ib' STARTING AT PAGE 3000000;
```

Using multiple disk controllers

If you have so much disk activity on multiple disks that you saturate the I/O bus, you should equip the server with multiple disk controllers, and connect the multiple drivers to the controllers as evenly as possible.

For example, if you have sixteen disk devices hosting database files, you might benefit from using four disk controllers, and attaching four disks to each controller.

Making drives specialized

A database server makes heavy use of both the operating system's virtual memory page file and of temporary disk space. If possible, equip the server with multiple disks and configure the virtual memory file, temporary directory, and database files on separate physical disk devices. This can use parallel I/O to the fullest advantage.

For example, on Windows NT you could locate the operating system files and *pagefile.sys* on *C:*, the temporary directory and infrequently-used files on *D:*, and database files on drives *E:* and higher.

Change the location of the virtual memory file with **Control Panel | System | Performance | Virtual Memory**.

Change the location of the InterBase temporary directory by either specifying a system environment variable INTERBASE_TMP, or editing the *ibconfig* file and specifying the path of the appropriate directory as a value for the TMP_DIRECTORY entry (see "Configuring sort files" on page 3-12).

Using high-bandwidth network systems

For client/server systems, hardware that supports high network bandwidth is as important as I/O capacity. The speed of the network often becomes a bottleneck for performance when many users are making demands on the network simultaneously.

Inexpensive 10 Base-T ethernet equipment is common today, but this technology is bare minimum for LAN configuration. It is recommended to use at least 100 Base-T for a high-performance network. The following graph illustrates relative bandwidth rates for various network interface technology.

Figure 11.2 Comparing bandwidth of network interfaces

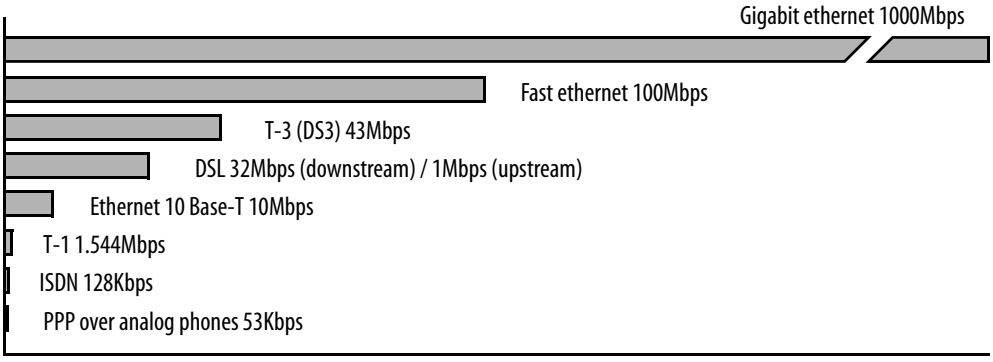

The maximum bandwidth of gigabit ethernet extends beyond the scale of the graph above.

At the time of this writing, most gigabit ethernet network interface cards (NICs) provide only 600 to 700Mbps bandwidth. Switches, routers, and repeaters also have constrained capacity. It is expected that the state of this technology will continue to improve.

It is recommended that you research reviews and experiment to learn the true throughput of all network hardware in your environment. The slowest component ultimately determines the true throughput.

Tip Network cables develop flaws surprisingly frequently. The result can be sporadic lost packets, for which operating systems compensate by automatically resending packets. This translates into mysterious network performance degradation. You should test network cables regularly. Replacing flawed cables is a low-cost way to keep your network running at peak efficiency.

Using high-performance bus

Bus is important for both I/O controllers and network interface hardware.

Figure 11.3 Comparing throughput of bus technologies

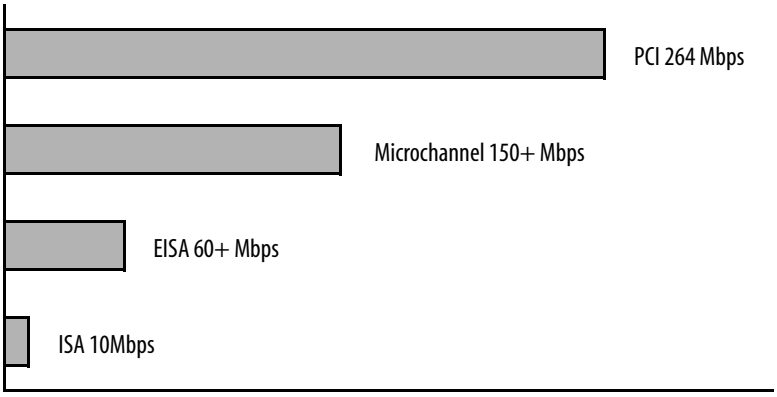

While 32-bit full-duplex PCI bus is capable of up to 264Mbps, PCI cards actually range from 40Mbps to 130Mbps.

Tip Use controllers on an integrated local PCI bus, it's faster than peripheral cards that plug into the motherboard.

Useful links

- The T10 Committee home page:
  `http://www.symbios.com/t10/`
  This is a useful place to find information on various storage interface technology.
- PC Guide Hard disk interface & configuration:
  `http://www.pcguide.com/ref/hdd/if/index.htm`
- The SCSI Trade Association:
  `http://www.scsita.org`
  News and vendor information about the state of SCSI technology and products.
- `The Gigabit Ethernet home page:`
  `http://www.gigabit-ethernet.org/`
- The Fibre Channel home page.
  `http://www.fibrechannel.com/`
  Fibre Channel (FC-AL) is an emerging extended bus technology for network, storage, video transmission, and clustering.

Operating system configuration

After you have equipped your server hardware appropriately, you should spend time tuning your operating system for server performance.

Disabling screen savers

Screen savers can have a serious impact on the performance of a server. Because servers are often set aside in a machine room, it's easy for the performance impact of a screen saver to be overlooked. Screen savers demand a surprising amount of CPU resources to run, and these programs run continuously, 24 hours a day.

Screen savers are evasive in their ability to disappear when a database administrator logs in to the console to diagnose a mysterious drop in performance. The server seems responsive to the database administrator as soon as she touches the server, but the speed degrades soon after she leaves the server.

Not all screen savers have the same performance cost. The Windows *OpenGL* screen savers perform continuous floating-point computations to draw three-dimensional shaded shapes in real time. They demand up to 90% of the system CPU, and cause InterBase and other services to slow to one-tenth their normal speed.

The Windows *Marquee* screen saver is one of the least demanding ones, especially when it is configured to pass text across the screen slowly. Some system administrators like to configure a Marquee on each screen in the machine room, to display the respective machine's hostname. This becomes a machine-name label, in raster form.

A screen saver can also be entertainment, but these should be reserved for workstations. A server in a machine room should be unattended, not used as a workstation.

If you must have phosphor burn protection for a monitor that you leave on, get an *Energy Star* approved monitor that has a power conservation mode. This mode blackens the screen after a configurable period of idleness. This not only protects against phosphor burn, but it conserves power. This is like a simple black screen saver, but it is handled by the electronics of the monitor, instead of by software.

The best option is to simply turn off the monitor when you aren't using it. This saves the phosphors, saves electricity, and decreases the amount of heat in the machine room.

Console logins

Don't leave the console logged in on a Windows database server. Even if the desktop is idle, it might use as much as 30% of the machine's CPU resources just maintaining the interface. You should log out of the server's console when you aren't using it. IBConsole enables you to perform most InterBase maintenance and monitoring tasks from another workstation, without logging in at the server's console.

Sizing a temporary directory

When you configure a temporary directory (see "Managing temporary files" on page 3-12), choose a location that has plenty of free disk space. For some operations such as building an index, InterBase can use a great deal of space for sorting. InterBase can even use an amount of space up to twice the size of your database.

The effects of insufficient temporary space include rapid virtual memory page faults, called *thrashing*, which causes a dramatic performance penalty. Another possible effect is a series of "I/O error" related messages printed to the *interbase.log* file on the server.

Use a dedicated server

Using a server for both workgroup file and print services and as a database server is like letting another user play a video game on your workstation. It detracts from the performance of the workstation, and it's not the intended use for the machine.

Use a secondary server as the file and print server, and a new machine for the database server. Alternately, use the secondary server for InterBase, depending on the relative priority of these tasks—the database server benefits from having a dedicated machine, even if it is not the fastest model available. Whatever is the most important service should be given the best machine as dedicated hardware.

If performance is a high priority, you can spend money more effectively by buying a dedicated machine instead of trying to increase resources such as RAM on a machine that is providing another competing service. Compare the cost of the hardware with the cost of having less than maximum performance.

Similarly, it is best to put a database on a dedicated drive, so that the database I/O doesn't compete with the operating system virtual memory paging file or other operating system I/O. See "Making drives specialized" on page 11-5.

Optimizing Windows NT for network applications

It is recommended to set the Windows NT server to optimize for network applications. Without this setting, you might see the CPU usage of InterBase peak for a few seconds every minute. With this setting, these peaks should vanish.

On Windows NT Server, the server is configured by default to give priority to filesharing services. You can change this configuration on the server: **Control Panel | Network | Services | Server**. In the Optimization panel, choose Optimize Throughput For Network Applications.

This change can result in a dramatic improvement of performance for InterBase, as well as other services.

Understanding Windows NT pitfalls

Windows NT has a peculiar way of balancing processes on SMP machines. If a process is exercising one CPU and the other CPU is relatively idle, Windows NT tries to switch the context of the process to the less burdened CPU. On a dedicated database server, the **ibserver** process is likely to be the only significant user of CPU resources. Unfortunately, Windows NT still tries to reassign the context of the process to the other CPU in this case. Once Windows NT has moved the **ibserver** process to the idle CPU, the first CPU becomes less burdened. Windows NT detects this and tries to move **ibserver** back to the first CPU. The second CPU becomes less burdened. This continues many times per minute, and the overhead of switching the process context between the CPUs degrades performance.

There are several possible solutions:

- Run **ibserver** on an SMP server that has enough other duties to occupy the other CPU
- Run **ibserver** only on a single-CPU machine
- Assign CPU affinity to the **ibserver** process:
  - **a** Launch the Task Manager
  - **b** Highlight the **ibserver** process
  - **c** Right-click to raise a window that includes CPU affinity settings

This technique works only if you run **ibserver** as an application, not as a service. If you run InterBase as a service, you must use the Windows API to programmatically set the CPU affinity of the **ibserver** process.

On some operating systems, using a ram disk is a technique for forcing very heavily used files to be in memory, but still allow them to be opened and closed like any other file. If you consider using a ram disk on Windows NT, be aware that the Microsoft ram disk utility for Windows NT uses paged memory to allocate the ram disk. The ram disk itself can be paged out of RAM and stored on the physical disk in *pagefile.sys*. It is futile to use a ram disk on Windows NT to create a high-performance file system.

Network configuration

This section describes performance considerations you should know when configuring a network configuration.

Choosing a network protocol

InterBase supports two protocols: TCP/IP when connecting to any server, and NetBEUI when connecting to a Windows server. See **"Network protocols"** on page 4-1 for more details.

NetBEUI

You can use NetBEUI on a network with fewer than 20 users without significant performance costs. Use TCP/IP if you have more active users on your network simultaneously.

NetBEUI is a network protocol designed for use on small local area networks. It is commonly used for filesharing services. It is a connectionless protocol, which means that it broadcasts packets to the entire network. This causes a growing amount of "noise" on a LAN. Noise, from the point of view of any given host, can be defined as network traffic that is not intended for that host. On a LAN with many hosts, enabling NetBEUI can overwhelm the network and reduce the available bandwidth for everyone to use. On most enterprise networks, IT experts discourage the use of NetBEUI.

TCP/IP

TCP/IP is a connection-based protocol, which means packets are routed to the intended recipient. This reduces the saturation of the network and the load on individual hosts. There is effectively more bandwidth available to all hosts, and a large number of hosts can share the same network with less performance penalty.

Configuring hostname lookups

Each host on a TCP/IP network has a designated IP address, and TCP/IP traffic is routed to hosts by address. TCP/IP requires a mechanism for clients to translate hostnames to their numeric addresses. Each client host can store the hostname/address associations in a file called *hosts*. You can alternately store this information on a central server, and the clients then retrieve the information on demand using a protocol called DNS. The client requests that the DNS server resolve a hostname, and the server returns the IP address. Then the client can use the IP address to communicate directly with the intended destination. In this configuration, the client must keep only one IP address locally: that of the DNS server host.

Depending on the load on the network and the DNS server itself, hostname resolution can take several seconds. This translates directly into delays when making a network connection. This is related to the message you might see in a web browser, "Looking up host name…" followed by, "Connecting to host name…" This indicates the delay while querying a DNS server to resolve a hostname.

You can speed up hostname resolution by adding the hostname/address mapping of the database server to the *hosts* file on the client computer. The client can resolve the hostname to its address much faster and more reliably by looking it up in a local file than by querying a service running on another host over the network. This reduces the hostname resolution delay when initiating connections to hosts listed in the local *hosts* file.

Note If you use this technique and later change the address of your database server, you must manually update the hosts files on each client workstation. Depending on the number of workstations in your enterprise, this can be tedious and time consuming. That's why DNS was invented, to centralize TCP/IP address administration. The suggestion to keep the database server address in a local file is intended to provide improved connection performance, but you should be aware of the administrative workload that it requires.

Tip If you object to the general IP address administration tasks required by using TCP/IP (independently from the DNS issue), consider using DHCP to simplify the task of assigning and tracking IP addresses of each host on the network. InterBase works in a DHCP environment as long as the client host has some means to resolve the server's IP address correctly at the time a client application requests an InterBase connection.

Database properties

Changing database properties can give an improvement in performance without changing anything in the design of your database. Applications require no change in their coding or design. Property changes are transparent to the client and database design.

Choosing a database page size

InterBase pages are 4KB by default. A typical production InterBase database gains 25 to 30 percent performance benefit by using this page size, relative to smaller page sizes. This page size results in better performance for the following reasons:

- Fewer record fragments are split across pages

  It is common for records to be larger than a single page. This means that InterBase fragments records and stores them on multiple pages. Querying a given record requires multiple page reads from the database.

  By increasing the size of a page, InterBase can reduce the number of multiple page reads and can store record fragments more contiguously.

- Index B-trees are more shallow

  Indexes are B-trees of pointers to data pages containing instances of specific indexed values. If the index B-tree is larger than one page, InterBase allocates additional database pages for the index tree. If the index pages are larger, InterBase needs fewer additional pages to store the pointers. It is easier for the database cache to store the entire B-tree in memory, and indexed lookups are much faster.

- I/O is more contiguous

It is fairly likely for a query to request successive records in a table. For example, this is done during a table scan, or query that returns or aggregates all records in a table. InterBase stores records on the first page that is unused, rather than ensuring that they are stored near each other in the file. Doing a table scan can potentially require retrieval of data by seeking all over the database. Seeks take time just as reading data takes time.

Any given page can store records from only one table. This indicates that a larger page is certain to contain more data from the same table, and therefore reading that page returns more relevant data.

- Default number of cache buffers is a larger amount of memory

  InterBase allocates the database cache in number of pages, rather than a fixed number of bytes. Therefore defining a larger page size increases the cache size. A larger cache is more likely to have a better hit rate than a smaller cache.

- Most operating systems perform low-level I/O in 4096 byte blocks

  InterBase performs a page read or write at the OS level by reading in 4096 byte increments regardless of the size of the database page. Therefore, by defining the database with a page size of 4096, the database I/O matches the low-level I/O and this results in greater efficiency when reading and writing pages.

Although 4KB seems to be the best page size for most databases, the optimal size depends on the structure of the specific metadata and the way in which applications access the data. For this reason, you should not consider the 4KB page size guideline to be a magic value. Instead, you should perform testing with your application and database under several different page sizes to analyze which configuration gives the best performance.

Setting the database page fill ratio

Data pages store multiple versions of data records, as applications update data. When a database is restored, the **gbak** utility fills pages with data only up to 80 percent of the capacity of each page, to leave space for new record version deltas to be stored, hopefully on the same page with the original record. But in a database that is used mostly for reading data rather than updating it, applications never benefit from this 80 percent fill ratio. In this case, it makes sense to restore data using the full capacity of each page. By storing 25 percent more data on each page, it reduces the amount of record fragmentation and increases the amount of data returned in each page read. You can specify the option to use all the space of every page for storing data during a database restore using the command:

```
gbak -c -use_all_space backup_file.ibk database_file.ib
```

Sizing database cache buffers

InterBase maintains a cache in the server's RAM of database pages currently in use. If you have a highly active database, you can gain some performance benefit by raising the default cache from its default of 2048 database pages. As with any cache system, at some point you find diminishing returns. Some experimentation with your particular application and database reveals that point.

See "Configuring the database cache" on page 6-24 for details about server cache settings.

The **ibserver** process running on an InterBase server maintains a cache in memory of recently used data and index pages. Like any cache, it depends on repeated use of data on a given page to help speed up subsequent access. In InterBase SuperServer implementations, the cache is shared by all clients connected to the database.

By default, InterBase allocates enough memory for 2048 pages per database. If the page size of the current database is 4KB, then **ibserver** uses 8MB of memory. If the page size is 8KB, then **ibserver** uses 16MB of RAM for cache. The InterBase API provides a method for any individual client to request that the size of the cache be higher. You can set a property on an individual database that establishes a different default cache size when any client connects to that database:

```
gfix -buffers 5000 database.ib
```

The default of 2048 assumes that the server has a sufficient memory configuration to allocate for 8MB of RAM per database. If memory is less plentiful on your server, or you have many databases that require simultaneous access, you might need to reduce the default number of cache buffers.

It is highly recommended to increase the cache size for a database if you have enough memory to accommodate it. Consider the following points:

- It is not useful to raise the cache size so high that the memory used by **ibserver** starts to page into virtual memory. That defeats the benefit of caching data from disk in memory.
- It is not useful to raise the cache size higher than the number of pages in the database (which you can view with View Database Statistics in IBConsole, or with the **gstat** command-line program). There's no benefit to this, since any given page from disk occupies only one page in the cache, and isn't duplicated.
- One block of memory is allocated for cache per database. If a client connects to two separate databases on one server, the **ibserver** process maintains two separate cache areas of memory. For example, if *database1.ib* has a default cache size of 8000 pages of 4KB each, and *database2.ib* has a default cache size of 10,000 pages of 2KB each, then while both databases have at least one connection, **ibserver** allocates a total of 32MB + 20MB of RAM.

You should experiment with larger cache sizes and analyze the performance improvements. At some point, you will observe diminishing returns. A typical application can achieve up to 30% performance increase from proper cache sizing.

The InterBase server does not use more than 512MB of cache per database, so you should not configure the number of cache buffers so high that it exceeds this amount of RAM.

Buffering database writes

InterBase on Windows platforms implements a *write-through* cache by default. Every write operation to a page in cache is immediately written out to the operating system's disk I/O, which itself might have a cache.

By contrast, a *write-back* cache defers flushing of the contents of a given cache page until a later time. InterBase performs multiple writes to a cache page in RAM before it writes the page out to disk. This results in better response time for the majority of write operations. Write-back cache consolidates I/O efficiently, and therefore it is much faster than write-through cache.

InterBase offers write-back cache as the default on UNIX and Linux, and as an option on Windows platforms. You can configure this at the database level using **gfix -write async** or by disabling forced writes for the database in IBConsole (**Database Properties | General tab | Options**).

The real benefit of using asynchronous writes (write-back cache) is about four times performance in the typical case. Some users have reported up to 20 times performance improvement from configuring asynchronous writes, in applications that make heavy use of write operations (INSERT, UPDATE, DELETE). The more writing an application does to the database—including write operations spawned by triggers—the more benefit the application gains.

The risk of asynchronous writes is that data in cache might be lost if the server has a power loss, or if **ibserver** exits abnormally for any reason. Write-through cache protects against data loss, at some performance cost. If you test your server host and client/server application thoroughly and they aren't susceptible to crashes, then it is highly recommended to use asynchronous writes.

Tip Use an uninterruptible power supply (UPS) to help protect your server against sudden power loss. A modest UPS is inexpensive relative to the cost of losing your data, and easy to install. This can allow you to gain the benefits of the asynchronous I/O mode in safety.

Database design principles

This section presents guidelines for database design techniques that benefit performance.

Defining indexes

Proper use of indexes is an important factor in database performance. Effective policies for defining and maintaining indexes can be the key to a very high performance client/server system. The self-tuning nature of indexes in InterBase greatly benefits performance, but you can gain some additional benefit by periodic maintenance tasks.

What is an index?

An index in InterBase is a Balanced-Tree data structure stored inside the database file that provides a quick lookup mechanism for the location of specific values in a table. Queries make use of appropriate indexes automatically by means of the cost-based optimizer, which analyzes the tables and columns used in a given query and chooses indexes that speed up the searching, sorting, or joining operations.

Defining indexes for some columns is part of designing a production database. Indexes dramatically improve performance of SELECT queries. The greater the number of rows in the table, the greater the benefit of using an index. Intelligently analyzing your database and defining indexes appropriately always improves performance.

Indexes incur a small cost to maintain the index B-tree data structure during INSERT and UPDATE operations. Because of this cost, it is not recommended to be overly liberal with index definitions. Don't create redundant indexes, and don't make an index on every column as a substitute for database usage analysis.

You shouldn't define an index for columns that have few distinct data values. For example, a column FISCAL_QUARTER might have only four distinct values over a potentially very large data set. An index doesn't provide much benefit for retrieval of data with this kind of distribution of values, and the work required to maintain the index tree might outweigh the benefits.

What queries use an index?

InterBase uses indexes to speed up data fetching for the following query elements:

- Primary and foreign keys
- Join keys
- Sort keys, including DISTINCT and GROUP BY
- Search criteria (WHERE)

In general, you should define indexes on all columns that you use in JOIN criteria or as sorting keys in an ORDER BY clause. You don't have to define indexes on primary or foreign key columns, because these table constraints implicitly create indexes.

What queries don't use indexes?

InterBase doesn't employ an index in the following operations, even if an index exists for the specified columns:

- Search criteria for CONTAINING, LIKE, and <> inequality operations
- Columns used in aggregate functions, like COUNT()
- Other expressions, like UPPER()

Directional indexes

Indexes are defined as either ASCENDING or DESCENDING. To sort in both directions, you need one index of each type. This is also very important if you are using a scrolling list in a Delphi form, or when using the *TTable.Last* method.

Normalizing databases

Design your database with proper normalization of data. Records that have lots of repeating groups of fields are larger than they need to be. Large records can increase the cost of sorting, and also cause records to span more pages than is necessary, resulting in more page fragmentation and needlessly large databases.

Denormalized table design can be more convenient for some types of client applications. You can use InterBase *views* and *stored procedures* to in effect store a denormalized query on the server, for convenient access from client applications. Meanwhile, the physical storage of the data is kept in a more efficient, normalized form.

See the *Data Definition Guide* for details on views and stored procedures.

Choosing Blob segment size

A Blob is a datatype with an unbounded size. It can be many megabytes in size, much larger than any database interface can handle in a single I/O transfer. Therefore, Blobs are defined as a series of segments of uniform size, and the I/O interface transfers Blobs one segment at a time.

Blobs are a special case because there is a special Blob page type, on which other datatypes cannot be stored. The data page for a record containing a Blob stores a Blob ID, which indicates which Blob page the Blob is stored on. A Blob is stored on the same page as the primary record version, if it fits. If it does not fit on that page, special pages are allocated for the Blob--as many as are required--and an index is stored on the primary page. Blob pages are never shared; either a Blob is on a normal data page, or it has a page to itself.

It is advantageous to define a Blob with a segment size equal to the page size. If both the page size and the Blob segment size are 4096 bytes, queries of large Blobs can achieve a data transfer rate of up to 20MB per second. InterBase ceases to be any kind of bottleneck in this situation; it is more likely that the hardware I/O bus, the network bandwidth, or the middleware are the limiting factors for throughput.

Database tuning tasks

This section describes ways you can perform periodic maintenance on your database to keep it running with the best performance.

Tuning indexes

Periodic maintenance of indexes can improve their performance benefit. You can write SQL scripts to automate these tasks. See **"Using SQL scripts"** on page 10-55.

Rebuilding indexes

Periodically, a B-tree data structure might become imbalanced, or it might have some values in the tree that have been deleted from the database (this should not happen in InterBase versions later than 5, due to index garbage collection).

You should periodically rebuild indexes by turning them off and on:

```
ALTER INDEX name INACTIVE;
ALTER INDEX name ACTIVE;
```

Recalculating index selectivity

The selectivity of an index is an indicator of its uniqueness. The optimizer uses selectivity in its cost-based analysis algorithm when deciding whether to use a given index in a query execution plan. If the selectivity is out of date and doesn't accurately represent the state of the index, the optimizer might use or discount the index inappropriately. This doesn't usually have a great performance penalty unless the selectivity is highly out of date.

You should recalculate the index selectivity if a change to the table affects the average distribution of data values:

```
SET STATISTICS INDEX name;
```

Performing regular backups

There are several performance-related benefits to doing periodic backup and restore of an InterBase database. See **"Benefits of backup and restore"** on page 8-1.

Increasing backup performance

- Disable garbage collection if you're just going to replace the database immediately anyway; this can make the backup execute faster.

- Back up to a different disk drive.

Increasing restore performance

- Restore from a different disk drive.
- Disable indexes on restore; this makes the restore execute faster so you have a usable database quickly. You must then have to manually activate the indexes after the restore is complete.

Tip Create a SQL script with all the ALTER INDEX statements necessary to activate your indexes, and keep that handy. Use it like a batch file with **isql -i script.sql** to help automate this procedure. You can create this script with this query:

```
SELECT 'ALTER INDEX ' || RDB$INDEX_NAME || ' ACTIVE;'
FROM RDB$INDICES
WHERE RDB$SYSTEM_FLAG = 0 OR RDB$SYSTEM_FLAG IS NULL;
```

You can get the database up and restored more quickly, then activate indexes afterwards. The data is accessible even if the indexes are inactive, but it's slower to query the tables.

Facilitating garbage collection

By default, InterBase databases have a built-in function to automatically sweep old record versions when they become too numerous. However, sweeping is partially inhibited by outstanding active transactions. If the server cannot do complete garbage collection, it has to do extra work to maintain each client's snapshot of the database.

Design your client applications to explicitly start and COMMIT transactions promptly, to reduce the number of outstanding transactions.

See "Overview of sweeping" on page 6-21 for more details on sweeping, garbage collection, and the database snapshot.

Application design techniques

This section describes general application programming methods for InterBase, that help to create high-performance clients.

Using transaction isolation modes

InterBase's multigenerational architecture requires that any query or other operation be associated with an active transaction. Without a transaction, an operation has no context with which to maintain its snapshot of the database. IBConsole and BDE tools do a certain amount of automatic transaction management, but it is helpful for performance to manually start and finish transactions.

In the InterBase server engine, a snapshot is generated by making a copy of the state of all other transactions in the database. This snapshot is static for the current transaction. This means that any data committed to the database after the snapshot is created is not visible to operations using that snapshot. This is the repeatable read transaction mode. Two identical queries made at different times are guaranteed to get the same result set, even if other clients are updating data in the database.

Starting a transaction and making a snapshot data structure for the new transaction incurs some amount of overhead. This overhead is magnified when using automatic transaction-handling, because the typical automatic transaction behavior is to start a new transaction and commit it for every statement executed against the database!

Another mode the default mode for BDE is called read committed. In this mode, the snapshot is updated every time the state of any transaction changes. This allows operations in the current transaction to view or act on data that has been committed since the snapshot was created. Updating the snapshot also costs a little bit in performance, so it is recommended to always use the repeatable read mode in InterBase. To do this, configure BDE driver flags to the value 512 or 4608.

Using correlated subqueries

Subqueries are SELECT statements which are included as a clause or expression within another statement. They are typically used to generate a value or result set that are used in conditions of the superior query.

A correlated subquery is one in which the conditions of the subquery are different for each row in the parent query, because they depend on values that vary from row to row. InterBase executes the subquery many times, once for each row in the parent query. Evaluating each row has a large cost in performance relative to a non-correlated subquery. InterBase optimizes non-correlated subqueries out of the loop, executes once, and uses the result as a fixed dataset.

Example as correlated subquery:

```
SELECT * FROM DEPARTMENT D
WHERE EXISTS (SELECT * FROM EMPLOYEE E
   WHERE E.EMP_NO = D.MNGR_NO AND E.JOB_COUNTRY = 'England')
```

Example as join:

```
SELECT D.*
FROM DEPARTMENT D JOIN EMPLOYEE E
   ON D.MNGR_NO = E.EMP_NO WHERE E.JOB_COUNTRY = 'England'
```

InterBase's optimizer executes a non-correlated subquery once, and uses the result set as many times as necessary in the parent query.

Sometimes a correlated subquery is necessary, given the semantics of the SQL language. However, these types of queries should be used with care and with the understanding that their performance is geometric in relation to the size of the dataset on which they operate.

Preparing parameterized queries

Any dynamic SQL (DSQL) statement must go through a cycle of parse, prepare, and execute. You can submit a DSQL statement to go through this process for each invocation, or you can separate the steps. If you have a situation where you execute the same statement multiple times, or the same form of statement with different parameters, you should explicitly prepare the statement once, then execute it as your looping action.

With *parameterized queries*, you can prepare a statement, but defer supplying the specific values for certain elements of the query.

InterBase supports parameterized queries in DSQL, for cases when a given statement is to be executed multiple times with different values. For example, loading a table with data might require a series of INSERT statements with values for each record inserted. Executing parameterized queries has a direct performance benefit, because the InterBase engine keeps the internal representation and optimization of the query after preparing it once.

Use parameterized DSQL queries in Delphi by following these steps:

1. Place a named parameter in the statement with the Delphi *:PARAMETER* syntax. in place of a constant value in a query. InterBase supports parameters in place constants. Tables and column names cannot be parameterized.
2. Prepare the statement. Use the TQuery method *Prepare*. Delphi automatically prepares a query if it is executed without first being prepared. After execution, Delphi unprepares the query. When a query will be executed a number of times, an application should always explicitly prepare the query to avoid multiple and unnecessary prepares and unprepares.
3. Specify parameters. For example, with the TQuery component, use the *ParamByName* method to supply values for each parameter in the query.
4. Execute the statement. SELECT statements should use the *Open* method of TQuery. INSERT, UPDATE, and DELETE statements should use the *ExecSQL* method. These methods prepares the statement in SQL property for execution if it has not already been prepared. To speed performance, an application should ordinarily call Prepare before calling *ExecSQL* for the first time.
5. Repeat steps 3 and 4 as needed.
6. Unprepare the query.

In some real-world cases involving repetitive operations, using parameterized queries has increased performance 100%.

Designing query optimization plans

The *optimization plan* describes the way the optimizer has chosen to execute a query. For certain types of queries, the optimizer might not select the truly optimal plan. A human can analyze different alternate plans and specify a plan overriding the optimizer's analysis. The result can be amazing improvements in performance for some types of queries. In some dramatic cases, this has been used to reduce a 15 minute query to three seconds.

The elements of plan selection are:

- Assigning indexes
- Combining indexes
- Determining join order
- Generating rivers
- Cost estimation
- Sort merges

InterBase supports syntax with the SELECT expression in embedded SQL and DSQL to allow the user to specify the PLAN for a query. The syntax also works with SELECT statements in the body of a view, a stored procedure, or a trigger.

It is beyond the scope of this chapter to describe in detail the syntax of the PLAN clause for specifying the execution plan, or techniques for analyzing queries manually. The section on SELECT in the *Language Reference* includes some examples of using PLAN.

Deferring index updates

Inserting and updating data requires indexes to be updated, which can cause performance to suffer during data INSERT or UPDATE. Some cost incurred while data is entered can result in a big performance win during later data queries.

To minimize the performance hit during INSERT, consider temporarily disabling indexes during high-volume INSERTs. This "turns off" the indexes, making them unavailable to help speed up queries, but also making them not be updated by data INSERTs. Then re-enable the indexes after INSERTing data. This updates and rebalances the indexes once for all the inserted data.

Application development tools

This section describes ways you can develop applications that are efficient, using various popular development environments and tools.

InterBase Express™ (IBX)

InterBase engineers at Borland have created a full-featured set of data-aware VCL components for use with the *TDataSet* architecture in Delphi 5. IBX can also be used with Borland's C++ Builder. See the *Developer's Guide* for full documentation of InterBase Express.

IB Objects

Another set of VCL components is available for projects with Delphi. It is designed to provide very sophisticated data component technology that is optimized for use with InterBase. The demo product can be downloaded from `http://www.ibobjects.com`.

Borland Database Engine

You should change the default values for BDE driver options in the BDE Administrator. This section provides guidelines for the driver options, and recommends values that you should use for better performance.

BDE driver flags

The recommended value for the DRIVER FLAGS is 4608.

By adding 512 to the DRIVER FLAGS in BDE Config tool, you specify that the default transaction mode is *repeatable read* transactions. This reduces the overhead that automatic transaction control incurs.

By adding 4096 to the DRIVER FLAGS, you specify that the InterBase SQL Links driver should use *soft commits*. Soft commits are a feature of InterBase that let the driver retain the cursor when committing changes. Soft commits improve performance on updates to large sets of data. When using hard commits, the BDE must refetch all the records in a dataset, even for a single record change. This is less expensive when using a desktop database, because the data is transferred in core memory. For a client/server database like InterBase, refreshing a dataset consumes the network bandwidth and degrades performance significantly. With soft commits, the client retains the cursor and doesn't perform a refetch.

Table 11.1 Matrix of BDE driver flags values

Driver flags	Isolation level	Commit type
0	Read committed	Hard commit
512	Repeatable read	Hard commit
4096	Read committed	Soft commit
4608	Repeatable read	Soft commit

Caveat: soft commits are never used in explicit transactions started by BDE client applications. This means that if you use explicit transaction start and commit, then the driver flag for soft commit is not used.

SQL passthru mode

The recommended value for this property is SHARED NOAUTOCOMMIT.

SQLPASSTHRU MODE specifies whether the BDE and passthrough SQL statements can share the same database connections. In most cases, SQLPASSTHRU MODE is set by default to SHARED AUTOCOMMIT. If however, you want to pass SQL transaction control statements to your server, you must use the SQL Explorer to set the BDE SQLPASSTHRU MODE to NOT SHARED. Depending on the quantity of data the client handles, you can achieve up to 10 times performance improvement by using the SHARED NOAUTOCOMMIT setting.

Use explicit transaction control and avoid autocommitted statements. Use the following methods: *TDatabase.StartTransaction*, and *TDatabase.Commit*.

SQL query mode

The recommended value for this property is SERVER.

The Active Server of InterBase includes a dynamic SQL parser and execution engine. In order for BDE to execute your SQL queries by sending them to the InterBase SQL engine, you must choose the value SERVER in this property. Otherwise, BDE parses and executes your query, which it does by fashioning a new SQL query and executing it by sending it to the InterBase server. There is no benefit to forcing BDE to reconstruct SQL that you have already written, only performance cost.

Visual components

This section describes visual components that developers commonly use in Delphi and C++Builder to access data from InterBase. Follow the recommendations below for better client/server performance.

Understanding fetch-all operations

In a client/server configuration, a "fetch-all" is the nadir of performance, because it forces BDE to request that the database generate a dataset again and send it over the network.

InterBase and most relational databases do not keep datasets in cache on the server in case the client requests a refresh. InterBase must execute the SQL query again when the BDE requests a refresh. If the query involves a large quantity of data, or complex joining or sorting operations, it is likely to take a long time to generate the dataset.

It is also costly for the server to transfer a large dataset across a network interface. It is more costly by far than it is for a desktop database like Paradox to return a dataset, because a desktop database typically runs locally to the application

It is often the case that software developers choose to use a relational database like InterBase because they are managing a larger amount of data than a desktop database like Paradox can handle efficiently. Naturally, larger datasets take more time to generate and to send over a network.

The person using the client application perceives that it has better performance if the user doesn't have to wait for refreshes. The less often the client application requests a refresh of the dataset, the better it is for the user.

Important A principle of client/server application design is therefore to reduce the number of costly refresh operations as much as possible.

TQuery

- *CachedUpdates* = False

Allows the server to handle updates, deletes, and conflicts.

- *RequestLive* = False

Setting *RequestLive* to False can prevent the VCL from keeping a client-side copy of rows; this has a benefit to performance because it reduces the network bandwidth requirement

- Below are some operations in which a TQuery perform a fetch-all. Avoid these as much as possible, or be aware of the cost of such operations.

Using the *Locate* method

You should use *Locate* only on local datasets.

Using the *RecordCount* property

It's convenient to get the information on how many records are in a dataset, but when using InterBase, calculation of the *RecordCount* itself forces a fetch-all. For this reason, referencing the *RecordCount* property takes as much time as fetching the entire result dataset of the query.

A common use of *RecordCount* is to determine if the result set of an opened TQuery contains any records, or if it contains zero records. If this is the case, you can determine this without performing a fetch-all by testing for both EOF and BOF states. If both end of file and beginning of file are true for the dataset, then no records are in the result set. These operations do not involve a fetch-all.

For example, for a given *TQuery* instance called *qryTest*:

```
qryTest.Open;
if qryTest.BOF and qryTest.EOF then begin
   // There are no result set records.
end
else begin
   // There are some result set records.
end;
```

Using the *Constraints* property

Let the server enforce the constraint.

Using the *Filter* property

For the TQuery to filter records, it must request a much larger dataset than that which it subsequently displays. The InterBase server can perform the filtering in a much more efficient manner before returning the filtered dataset. You should use a WHERE clause in your SQL query. Even if you use a WHERE clause, any use of the *TQuery.Filter* property still forces a fetchall.

TTable

The *TTable* component is designed for use on relatively small tables in a local database, accessed in core memory. *TTable* gathers information about the metadata of the table, and tries to maintain a cache of the dataset in memory. *TTable* refreshes its client-side copy of data when you issue the *TTable.post* method and when you use the *TDatabase.rollback* method. This incurs a huge network overhead for client/server databases, which tend to have larger datasets and are accessed over a network. You can observe the activity of *TTable* with the SQL Monitor tool. This reports all calls to the BDE and InterBase API.

Though *TTable* is very convenient for its RAD methods and its abstract data-aware model, you should use it sparingly with InterBase or any other client/server database. *TTable* was not designed to be used for client/server applications.

InterBase Limits

This appendix defines the limits of a number of InterBase characteristics. The values the following table lists are design limits, and in most cases are further restricted by finite resource restrictions in the operating system or computer hardware.

Table 11.2 InterBase specifications

Characteristic	Value
Maximum number of clients connected to one server	There is no single number for the maximum number of clients the InterBase server can serve—it depends on a combination of factors including capability of the operating system, limitations of the hardware, and the demands that each client puts on the server. Assuming a "normal" type of client application that executes database operations from human interaction, and a modern server platform (Pentium 150MHz+, 64MB RAM), expect the InterBase server to comfortably handle up to 150 clients. This is a guideline, not a guarantee. Applications that engage in high levels of contention or that perform complex or high-volume operations could cause the practical number of clients to be fewer. Note also that some operating systems do not have the technology to serve 150 incoming network connections.
Maximum database size	The maximum addressable file size for a single file is 2GB on Windows 98, 4GB on most other platforms. Refer to your operating system documentation to verify file size limits. Combined with the multifile database feature of InterBase, this allows many terabytes of addressable file space.
Maximum number of files per database	By design, 2^{16} (65,536), because the files are enumerated with an unsigned 16-bit integer. Shadow files count toward this limit. This is a design parameter of InterBase, but most operating systems have a much lower limit on the number of files that a single process can have open simultaneously. In some cases, the OS provides a means to raise this limit. Refer to your OS documentation for the default open files limit, and the means to raise this limit.
Maximum number of cache pages per database	65,536; for the sake of performance, a more practical upper limit would be 10,000. Total size of cache pages should never exceed 50% of memory.
Maximum number of databases open in one transaction	No restriction. The parameters in a transaction parameter buffer comprise a linked list, so there is no limit except that imposed by system resources.

Table 11.2 InterBase specifications (*continued*)

Characteristic	Value
Maximum number of tables per database	By design, 2^{16} (65,536), because tables are enumerated with a 16-bit unsigned integer.
Maximum versions per table	255; then no more metadata changes until the database has been backed up and restored.
Maximum row size	64KB. Each Blob and array contributes eight bytes to this limit in the form of their Blob handle. Systems tables (tables maintained by the InterBase engine for system data) have a row size limit of 128KB.
Maximum number of rows and columns per table	By design, 2^{32} rows, because rows are enumerated with a 32-bit unsigned integer per table. Number of columns in a row depends on datatypes used. One row can be 64K. For example, you can define 16,384 columns of type INTEGER (four bytes each) in one table.
Maximum number of indexes per table	By design, 2^{16} (65,536), because indexes per table are enumerated with a 16-bit unsigned integer.
Maximum number of indexes per database	By design, 2^{32}, because you can create 2^{16} tables per database, and each table can have 2^{16} indexes.
Maximum index key size	Starts at 256 bytes for a single-column key, and 200 for multicolumn keys; subtract four bytes for each additional column. Example: a single-column CHAR key can be up to 256 – 4 = 252 bytes; a three-column key must add up to 200 – 12 = 188 bytes. Note that multibyte character sets must fit within the key by counting bytes, not by counting characters. For example, a single-column key using 3-byte UNICODE_FSS characters can have a maximum of (256 – 4) / 3 = 84 characters.
Maximum number of events per stored procedure	No restriction by design, but there is a practical limit, given that there is a limit on the length of code in a stored procedure or trigger (see below).
Maximum stored procedure or trigger code size	48KB of BLR, the bytecode language compiled from stored procedure or trigger language.
Maximum Blob size	The size of the largest single Blob depends on the database page size: 1KB page size: largest Blob is 64MB 2KB page size: largest Blob is 512MB 4KB page size: largest Blob is 4GB 8KB page size: largest Blob is 32GB A Blob is a stream of many segments. The maximum Blob segment size is 64KB.

Table 11.2 InterBase specifications (*continued*)

Characteristic	Value
Maximum tables in a JOIN	No restriction by design, but the task of joining tables is exponential in relation to number of tables in the join. The largest practical number of tables in a JOIN is about 16, but experiment with your application and a realistic volume of data to find the most complex join that has acceptable performance.
Maximum levels of nested queries	There is no restriction by design. The practical limit depends on the type of queries you nest. Experiment with your application and a realistic volume of data to find the deepest nested query that has acceptable performance.
Maximum number of columns per one composite index	16.
Levels of nesting per stored procedure or trigger	750 on Windows platforms. 1000 for UNIX platforms.
Maximum size of key in SORT clause	32KB.
Range of date values	January 1, 100 a.d. to February 29, 32768 a.d.

Chapter

12

Data Replication

This chapter documents Synectics Software's InterBase replication server, IBReplicator. It covers the following topics:

- IBReplicator and its components
- Data replication
- How to use IBReplicator

About IBReplicator

IBReplicator is a third-party data replication system created by Synectics Software Ltd. You can use IBReplicator to ensure that changes to the data in any InterBase database can be duplicated in any number of other InterBase databases, even databases with different structures. IBReplicator is included on the "Server" version of the InterBase CD-ROM. If you do not have this available, but are interested in IBReplicator functionality, you can download an evaluation copy of IBReplicator from Synectics Software, at **http://www.synectics.co.za/**.

IBReplicator consists of several executables, in two categories:

- Configuration tools
- Replication server

Configuration tools

Below are the available IBReplicator configuration tools:

Table 12.1 IBReplicator configuration tools

Executable	Description
IBRplManager.exe	Replication Manager: a GUI tool that runs only on Windows operating systems, but can be used to configure replication for all platforms.
IBScheduler.exe	Replication Scheduler: a Windows NT/2000 service, used to schedule replication on any platform
IBReplSrvcInstall.exe	A GUI tool, used to install the Windows NT/2000 Service versions of various InterBase/Replication tools

Replication server

On Windows platforms, there are two available executables:

- **IBReplserver.exe** is the GUI version used for testing and replication on Windows 98/ME
- **IBRepl.exe** is the Windows NT/2000 Service version

On UNIX and Linux platforms, there are two available executables:

- **replserver** is the replication server; it runs as a daemon
- **replmgr** is the command-line tool used to start and stop the replication server

InterBase support

The Replication Server can replicate from/to InterBase 6 databases, or from/to InterBase 5 databases, but the Replication Server and the Replication Manager must run on a machine that has an InterBase 6 client installed. The Replication Configuration database must be created on an InterBase 6 server.

Requirements

This section describes supported platforms, supported InterBase versions, and OS requirements.

Windows system requirements

Replication Manager, Replication Server

Microsoft Windows NT4.0, Windows 2000, Win 98

Memory: 16 megabytes minimum; 64 or more recommended

Processor/Hardware model: 486 minimum; Pentium II recommended

Disk space: 2MB for each application

InterBase Version: 5.*x* and 6.0

Sun Solaris 2.5.x or 2.6.x system requirements

Memory: 32 megabytes minimum; 64 or more recommended

Processor/Hardware model: SPARC or UltraSPARC

Disk space: 2MB for the application

InterBase Version: 5.*x* and 6.0

Linux system requirements

Essentially the same as for Windows NT/2000. The replication server runs on any version of Linux supported by InterBase.

IBReplicator features

This section describes the range of functionality available in IBReplicator:

Overview

- **Fast** Replication occurs directly between servers: there are no intervening layers of processing from database engines or drivers, for example. Our benchmarks indicate that actual replication speeds range from five operations per second over a 28.8 dial-up connection to 200 operations per second between two 200Mhz Pentium machines on a 10BaseT network, where an "operation" is an insert, an update or a delete.
- **Small** IBReplicator consists only of the necessary code; there is none of the overhead associated with any form of middleware. The Replication Engine requires 2MB of disk space on the Server and a further 2MB on the Replication Manager machine. In a Windows environment these may run on the same machine.

InterBase advanced features

- **Advanced datatypes** IBReplicator can replicate all supported InterBase datatypes, including Blobs and Arrays. It handles multi-segment primary keys where each segment can be any supported InterBase datatype.
- **Event alerters** Replication can occur in response to database events.
- **Internationalization** IBReplicator international character sets are supported.

Ease of use

A replication tool is, in its nature, a complex piece of software, especially if it is a highly configurable one, like IBReplicator. Nonetheless, Synectics Software's Replication Server is exceptionally easy to install and use:

- **Point-and-click configuration** This tool allows you to select which tables and fields are to be replicated, and to view and edit optional settings; it also generates the required triggers on the source database for you. A Publish and Subscribe model is used to configure replication schemas.
- **Minimal configuration** Target databases need no configuration at all. The Replication Server can automatically map all tables and fields in the source database to tables and fields of the same name on the target database.
- **Manual Conflict Resolution** This tool allows logged conflicts to be resolved manually.
- **Help** The Replication Manager includes a comprehensive, context-sensitive online help file with detailed explanations of procedures and controls.

Components of IBReplicator

IBReplicator is a system that includes several components, of which the management tool is the most used. The following is an overview of these components.

- **The Replication Server** performs the actual replication of your data. It is available as a service under Windows NT/2000 (**IBRepl.exe**).

Note On Windows NT/2000, the utility (*IBReplSrvcInstall.exe*) installs, runs, stops, and removes the Replication service. Make sure the correct path to *ibrepl.exe* is displayed when you install the service.

- **The configuration database** saves the details of which data to replicate, and where it should be replicated to (these replication specifications are called *schemata*). You can create as many configuration databases as are required.
- **The Replication Manager** (*IBRplManager.exe*) provides a central user interface from which you can define schemata, manage your configuration databases, and run several supplied utilities. It also includes a comprehensive, context-sensitive online help file with detailed explanations of procedures and controls.
- **The scheduler** defines the intervals at which replication should occur. This tool is available as a service under Windows NT/2000 (*IBScheduler.exe*), and also as an application so that it can be run under Windows 98/ME (*IBReplScheduleManager.exe*). It is run from within the Replication manager.
- **Replication Monitor** This tool provides a real-time graph showing the status of ongoing replications. It also provides information and statistics on all related connections, throughput and activity in the environment. The Monitor can be found in the Tools menu within the Replication Manager.

About data replication

Essentially, a replication server ensures that changes to a *source* database are duplicated on other, *target*, databases so that all databases contain the same data. These source and target databases can be on the same machine, but are usually separated and need to be connected with a local- or wide-area network, or with a dial-up connection.

In principle, this is a straightforward idea, but reality quickly introduces complications. For example:

- **N-way replication** A single database can be *both* a source of data *and* a target of other databases at the same time (see "Choosing the source database" on page 12-18).
- **Complex replication** Source and target databases are often structured differently: tables can have a different number of columns with different names and data-types, some tables can be present in some databases and absent in others, and so on (see "Choosing replicated tables and stored procedures" on page 12-21).
- **Conflicts** Typically, target databases are maintained by their own users, who are adding, deleting and updating rows without knowing about the changes made on a distant source database at the same time. So, a replication server might discover that a row that has been updated is entirely missing on a target, or it might find it needs to add a row to a target already using its primary key (see "Customizing default settings" on page 12-16 and "Choosing the source database" on page 12-18).
- **Replication timing** In some contexts, it is enough for replication to happen every so often, perhaps monthly, or daily, or only in response to explicit requests. This *asynchronous* replication requires only a low-end network setup and can be scheduled to occur during idle times such as at night or over weekends, but, of course, the various databases are almost always more or less out of date. In contexts where databases must always be up to date, *synchronous* replication is needed. In this case, changes are replicated as they occur, with all the consequent pressures on the network. In general, a compromise solution is adequate, with critical data being replicated near-synchronously, and less topical data being replicated asynchronously (see "Customizing default settings" on page 12-16 and "Choosing the source database" on page 12-18).
- **Data subsetting** Often, only certain rows in a table should be replicated; in some cases it is necessary to replicate certain rows to one target, and other rows to other databases (see "Choosing replicated tables and stored procedures" on page 12-21).
- **Synchronization** In general, changes in source databases only need to be replicated. You don't need to synchronize. Synchronization is useful for tasks such as filling a new, empty database with all the data already in a source database and for bringing a database restored from an old backup up to date. Note that

synchronization is used only under these fairly rare circumstances. If a target database has simply been offline, there is no need to synchronize it as normal replication will catch up with the backlog.

IBReplicator handles all these cases, and it takes pains to make the configuration of a replication environment as straightforward and yet as flexible as possible.

You will appreciate, however, that it is imperative that you *design* your environment carefully. For example, it would be a mistake to replicate data from database A to database B, and then to cause B's changes to be replicated back to A, and this can happen easily in cases where A replicates to B, C, D and E, while B is busy replicating to X, Y and Z, which in their turn are replicating backwards and forwards, and so on.

Note Some authorities discuss "published" databases replicating to "subscribed" databases. They are using a metaphor which sees a database as a sort of journal which provides data to other databases that are on its mailing list. IBReplicator simply sees a database as being either a source of information, or as a target for it (or both, of course).

These two models are not quite identical. The published/subscribed metaphor implies that the published side of the replication needs to be configured to provide the data for all its subscribers. This complicates complex replication situations where different subscribers need different subsets of the available data, and also involves reconfiguring the publisher whenever a subscriber with different needs is added.

IBReplicator's source/target model allows each target to specify whatever data it needs, but it does mean that identical targets need to be configured one by one (see "Cloning targets" on page 12-21).

How IBReplicator works

As you define your replication schemata, the Replication Manager maintains the configuration database in which they are stored. When a schema has been defined, you complete its implementation by telling the Manager to create IBReplicator's *system objects* in the source database.

These system objects are triggers and tables that ensure that every change to the data in the source database is saved in a table called the Replication Log that stores the action that was performed (an insert, update or delete) and the primary key value of the row affected.

The basic sequence

The Replication Server can receive an instruction to replicate from any of several sources:

- Its own internal timer

- Your explicit request
- A timed signal from the **Scheduler**
- An InterBase event

When the server receives such an instruction, it first consults the configuration database to determine which databases are involved. Next it queries its system replication log table in the source database to find the rows that have changed. Finally, it retrieves those rows (which contain the very latest data), and duplicates the action in each target database:

- An *insert* is replicated by inserting a new row with the same primary key values and the same values in each column that is to be replicated.
- An *update* is replicated by finding the row with same primary key and updating its replicated columns to their new values.
- A *delete* is replicated by finding the row with the same primary key and deleting it.

This design helps make IBReplicator easy to use because you need to configure only your source databases; their targets are left untouched.

Furthermore, the design makes the replicator flexible enough for you to organize your network in any way that suits your needs. On the one hand, even the most complex of configurations can be implemented with only a single instance of IBReplicator running on a single server with a single configuration database. But you can also have the replicator running on several or all of your servers, and each instance of the replicator can use its own configuration database (or databases), or various instances can share a configuration, all at your convenience.

IBReplicator's strategies

- **Row-level replication** Also called Data Subsetting. This enables the limiting of data sent to a particular target to a particular domain. For example only replicate rows for Departments 'A' and 'C' to Machine X, and replicate rows for Departments 'S' and 'Q' to Machine Y.
- **Sequence of events** All changes are processed in exactly the same sequence on the target database, as they were originally done on the source database.
- **Sequence of changes** Changes are replicated in their original order.
- **Multisegment primary keys** IBReplicator recognizes primary keys made up of multiple columns, and allows each column to be of any InterBase datatype.
- **Supertransactions** IBReplicator replicates only complete transactions. The Replication Server bundles transactions into supertransactions for efficiency.
- **Encryption** All passwords that are stored in a database or in the system registry are encrypted.

- **Field names** Target tables need not have exactly the same field names and datatypes as their sources.
- **Altered replication schemas** The server can reload configuration settings on command from the Configuration Tool.

Resolving precedence issues

Each Source/Target database pair can have its own conflict resolution settings. IBReplicator provides three ways to handle cases where replicated data conflicts with existing data in the target database:

- Priority-based

Databases can be given priorities, and the database with the higher priority takes precedence.

If the source database has precedence the following occurs:

- An UPDATE finding no identical key record in the target database is converted into an INSERT.
- An INSERT finding a record in the target database with an identical key is converted into an UPDATE.
- A DELETE finding no identical key record in the target database is ignored.

- Time-stamped

Older rows in the source database will not overwrite newer rows in the target.

- Master/slave

The source database always takes precedence.

Replication involving newer InterBase datatypes

Beginning with InterBase 6, data in NUMERIC and DECIMAL fields that have a precision greater than 9 is stored as INT64. This means that when replicating to InterBase 5, the data may not fit, especially if the InterBase 5 field is being stored as an INT32. The Replication Server checks to see if the data fits, and returns an error if it does not.

Current InterBase TIME and DATE and TIMESTAMP fields replicate to InterBase 5 DATE fields. IBReplicator adds the current date to the TIME field in order to make it into a proper InterBase 5 DATE field.

Operation logging

Each Source/Target database pair can have its own error/information logging settings and log file.

IBReplicator can record its transactions in a log that can be either a window on screen or a disk file according to a minimum severity level. You can choose to have it write any of the following items to the error log:

- All errors (for example: can't connect to target, record already exists, no record to delete, record locked).
- Replication statistics: performance and warnings
- Database connections made
- Values of Key Fields
- SQL Statements generated

Viewing schema

IBReplicator provides a *Schema view* utility that provides a visual representation of your replication setup. It can help you to confirm that your design has been implemented correctly.

To run the Schema View utility, go to the Replication Manager and choose **Tools | Schema view**.

Running the Replication Server

The server portion of IBReplicator can run on Windows, Linux, and Solaris. On Windows it can run as either an application or a service. This section describes the binaries and their functions.

Windows platforms

There are two executables for the various windows platforms *IBReplserver.exe*, which is a standard GUI executable, and IBRepl.exe which runs as a service. The service version can only run on Windows NT or on Windows 2000, while the GUI version can run on any of the windows platforms.

A utility, *IBReplSrvcInstall.exe*, has been provided to install/run/stop/remove the Replication service. Make sure that the correct path to the *Ibrepl.exe* executable is displayed when installing the service.

The Replication Manager contains a Replication Scheduler, which can be configured to invoke the Replication Server at the required intervals. This scheduler can also be run as a service. The required executable is called **IBScheduler.exe** and is installed and removed the same way as the Replication Service.

UNIX and Linux platforms

Two executables are supplied on these platforms:

replserver The binary for the server itself

replmgr The binary for the utility that starts and stops the Replication Server

- To start the **replserver** executable directly, pass the relevant parameters to either **replserver** or **replmgr**.
- To stop the server, you must use **replmgr**.

Replication Server parameters:

```
replserver -u[ser] <username> -p[assword] <password>
   -r[ole] <rolename> -mode <window/daemon> configdbpath
```

Replication Manager parameters:

```
replmgr -u[ser] username -p[assword] password
   -r[ole] <rolename> -a[ction] <start/stop/force>
   -mode <window/daemon> configdbpath
```

Notes:

1. You should not run the Replication Server more than once against the same Configuration database. The results are unpredictable, since the same data may be replicated more than once.
2. It is permissible to run the server more than once on the same machine, as long as each instance of the Replication Server is run against a different configuration database.
3. When stopping the Replication Server with '-action stop', you need to specify the same config database path as when you started the server. This is because replmgr uses InterBase Events to communicate with the running Replication Server, so it must connect to the same database.
4. If you are running multiple instances of the Replication Server on one machine (each connecting to its own *config* database), you will need to run '**replmgr -action stop**' multiple times, once each for each active configuration database.
5. There are four available environment variables:

Variable	Use
ISC_USERNAME	Contains the required username
ISC_PASSWORD	Contains the required password
ISC_ROLENAME	Contains the required role name
IBPATH	Contains the name of the directory where the server executable is located if it is not in */usr/interbase/bin*

Using IBReplicator

Once you have installed IBReplicator, you need to define your replication schemata so that the Replication Server can identify your source and target databases, log into them, and replicate only the desired data. You use the management tool, Replication Manager, to design replication schemata.

- The replication engine doesn't have to be running on either the Source or the Target database server for replication to occur; it can be running on any machine on your network or intranet.
- The Replication Server is available for most platforms supported by InterBase, including flavors of UNIX, but the management tool is a Win32 application: it must be run on a Windows machine that has access to the machine running the Replication Server.
- IBReplicator creates a default log file in which all fatal errors, including those that occur during start-up, are logged. This file is stored in InterBase's root directory on the machine where the replication server is running; it is called *Replication.log*. You can specify your own log file when you define your replications later on.
- You can extend the standard replication capabilities by writing additional replication logic in stored procedures.

Using Replication Manager

You use Replication Manager to design replication *schemata*: replication specifications.

Replication Manager's main window provides three important controls, which correspond to the three steps needed to define a replication schema for IBReplicator to implement for you:

1. *Configurations* combo-box, where you choose the configuration database that you need to work on.
2. *Databases* tab, where you identify the databases to be involved in replication.
3. *Replications* tab, where you to define the replication schemata which identify the databases, tables, columns and rows to be replicated.

Each of these steps is discussed in the following sections of this chapter.

Important The Replication Manager includes a comprehensive, context-sensitive online help file with detailed explanations of procedures and controls.

Managing configuration databases

When you run Replication Manager for the first time, it notifies you that you must create a configuration database. The Replication Server uses this database to establish what it is supposed to replicate, where it should find the data, and where that data should be sent.

Configuration databases are defined in a special Configuration Databases dialog that you access from *File* | *Configurations* in the Replication Manager. With this tool, you can:

- Create and drop configuration databases.
- Edit each configuration database's connection parameters.
- Make a configuration database the default.
- Add existing configuration databases to the list of those known to IBReplicator, and remove them from the list without dropping the database.

All these functions are available from the toolbar, and from the **Configurations** menu.

Creating a configuration database

To create a new configuration database, follow these steps:

1 In Replication Manager, choose **File** | **Configurations** to display the Configuration Databases Editor.

Figure 12.1 The Configuration Editor

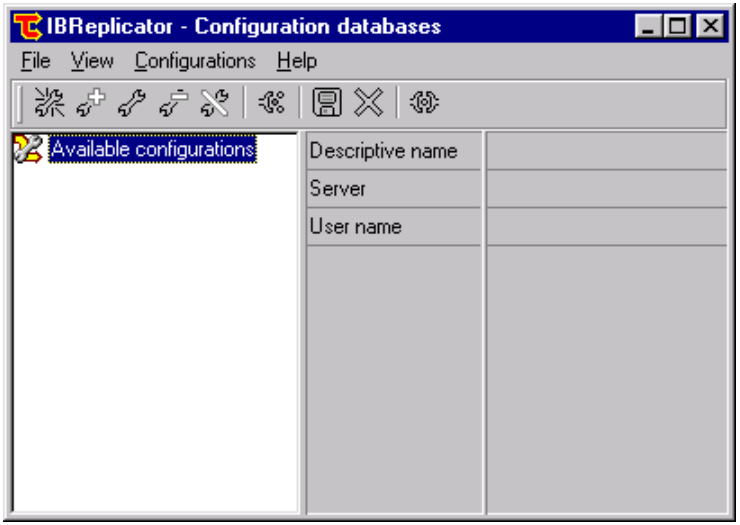

2 In the Configuration Databases Editor, choose *Configurations* | *Create*.

3 Give the new configuration a descriptive name. This name allows you to identify your configurations when you select the configuration database to edit in the management tool, so the name should be self-explanatory and unique.

4 Supply the path and filename for the new database, and the user name that the Replication Server should use when connecting to it.

5 Choose **Configurations | Save** to create the new database.

Working with configurations

- **To add a configuration database**, choose **Configurations | Add** and follow the same steps used to create a new configuration (see above). The one difference is that the filename you supply *must* identify a configuration database that has already been created.

Each instance of the Replication Server must have its own configuration database; instances cannot share a configuration.

- **To remove a configuration database**, choose **Configurations | Remove**. Removing a configuration does *not* destroy the database; it simply means that IBReplicator can no longer access it.
- **To destroy a configuration database**, choose **Configurations | Destroy**. This drops the database. Be aware that a dropped database is deleted from the disk!

The default configuration

One of your configuration databases must be identified as the default:

- **To make a configuration the default**, choose **Configurations | Default**. At the prompt, supply the password that the replicator and its tools should use when connecting to that database.

When you run the Replication Server as a *service* under Windows NT/2000 (see "About IBReplicator" on page 12-1), it always connects to the default configuration. However, when you run the server as an *application* by running **IBReplicator.exe**, you can specify which configuration to use.

Registering databases

When you have chosen or created a configuration database, you are ready to define the replications that you need, and this is done in Replication Manager's main window. The first step is to identify the databases to be involved using the Databases tab.

You need to register all the databases involved in the replication, both those that serve as sources of data, and those that receive it (the targets). *Registering* a database identifies a database for IBReplicator, and provides values for the parameters it needs to connect to that database.

Important Registering a database does *not* create the database, and does not cause any replications to happen.

Figure 12.2 Replication Manager database tab

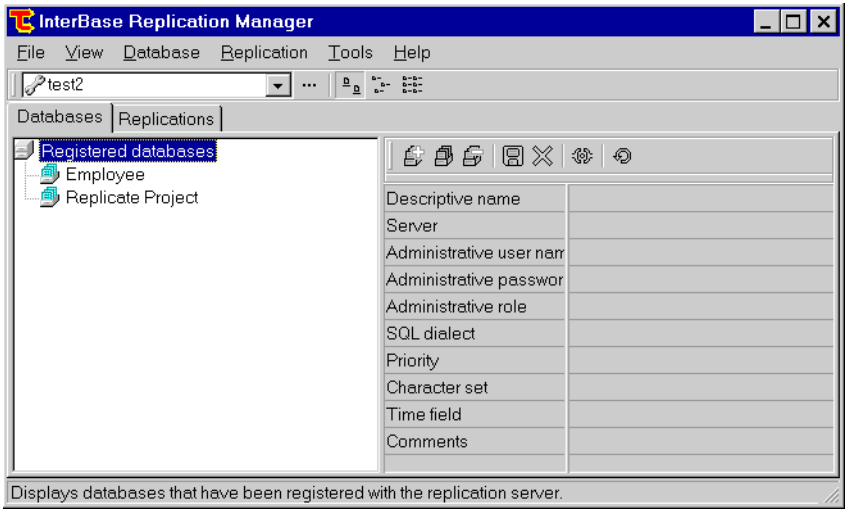

The Databases tab provides three controls:

1. A Tree view that lists all the databases that have been registered and saved in the current configuration database.
2. A Field Editor that lists the connection parameters, displays their current values, and allows you to edit those values.
3. A Toolbar that gives quick access to the commands relevant to database registration. These commands are also available from the Databases menu.

Some of the database parameters *must* have values; the others can be set if they are applicable, and ignored if they are not. The essential parameters are:

- *Server* The network path and filename of the database file. Make sure the syntax of this path conforms to the syntax for the kind of network connection that InterBase uses to connect to that database.
- *Administrative Username* and *Password*. The username and password which the management tool should use when connecting to the database. Compare the *replication* username and password used by the Replication Server when replicating (see **"Choosing the source database"** on page 12-18 and **"Choosing the target databases"** on page 12-20). Of course, the administrative and replication usernames and passwords can be identical.
- The *descriptive name* is another optional but important field. It is optional in the sense that the Replication Server doesn't use it. It is important because it makes your replication schemata easy for users to follow. A name of "Head Office (*employee.gdb*)" is more helpful than the default "New database".

Defining replication schemata

When you have registered the databases that will be involved in replication, your next and final step is to identify the data to be replicated. This is done by defining replication *schemata* on the Replications tab of Replication Manager's main window. This tab contains two controls:

1. A Tree view that displays all the objects to be involved in replications, including source and target databases, tables, and columns.
2. A control panel that presents commands applicable to the selected node in the Tree view. These commands are all available from the Replication menu as well, but only some of the available commands are relevant at any time, and just which these are depends on the node selected in the Tree view. For example, if you have selected a target database, commands for identifying replicated columns are not applicable because you have to identify the table that contains those columns first.

Defining a replication

To specify which data should be replicated and where it should be replicated to, on the Replications tab follow these steps:

Step	Task	Node	See page
1 (optional)	Customize default settings	Defined Replication Schemata	16
2	Create schema	Defined Replication Schemata	17
3	Choose source database	Source Database	18
4	Choose target database(s)	Target Databases	20
5	Choose replicated tables	Replicated Tables	21
6	Choose replicated columns	Data Columns	26
7	Create system objects	A source database node	27

Figure 12.3 Replication Manager Replications tab

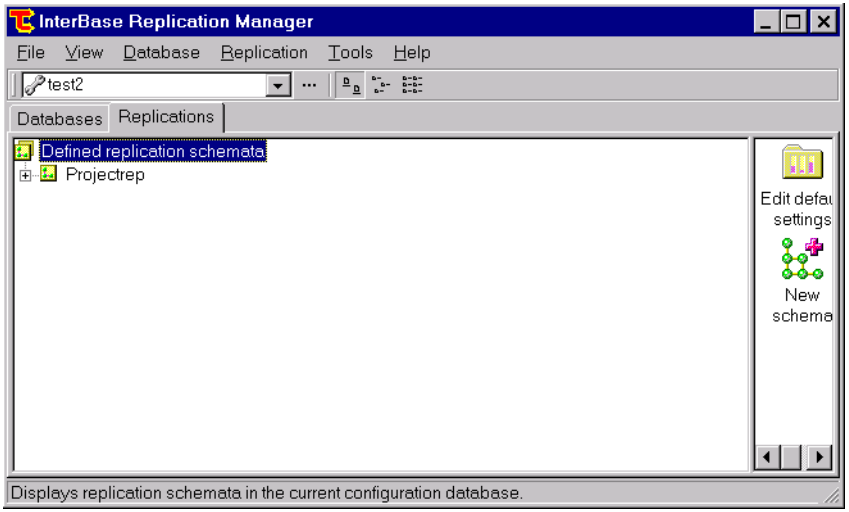

Customizing default settings

Every schema you define has a variety of customizable settings that control replication intervals, conflict resolution strategies, synchronization, and event logging. All of these have been initialized to reasonable values for you, and it is these values that will be supplied to each new schema you create.

To customize IBReplicator's default settings, and to define a new schema, select the root node, Defined Replication Schemata, in the replication Tree view.

To change the default values for these settings, double-click the Edit Default Settings icon in the replication list view or choose **Replication | Default settings** to open the Default Settings dialog.

Setting	Description
Replication timing	Replication can occur at scheduled intervals (defined with the scheduler utility), or in response to an event alert from a source database.
Conflict resolution strategy	When new data from a source database has a primary key that is already used by one of the target databases, IBReplicator can: • Preserve the data from the database with higher priority (as defined when the databases were registered), or • Preserve the most recent version of the data, or • Replicate the source database's version regardless. You can have the Replication Server log conflicts for you to resolve later by choosing **Tools | Conflicts** from Replication Manager's main menu.
Synchronization	*Synchronization* involves updating a target database that is missing data that should have been replicated to it. You can enable synchronization, and you can allow *reverse* synchronization, where the target updates its source.
Event logging	The Replication Server can log its activities for you to inspect, either to a window or to a file.

These defaults are stored in the configuration database, so different configurations can have different defaults. Changing default settings does not change the settings of existing schemata.

Creating schemata

A replication schema contains information about which data should be replicated, where the data can be found, and where it should be sent.

To define a new schema

1. Select the replication Tree view's root node.
2. Run the New Schema command by either double-clicking its icon in Replication Manager's control panel, or by choosing **Replication | Schema | New** from the menu. This opens the New Schema dialog.
3. Give your schema a descriptive name.
4. If desired, you can also override the default synchronization settings
5. Choose OK to add a schema node to the Tree view as a child of the Defined Replication Schemata node.

To edit a schema

To edit a schema's name and synchronization settings, select a schema node and double-click Edit Schema or choose **Replication | Schema | Edit**.

To delete a schema

To delete a schema, double-click Delete Schema or choose **Replication | Schema | Delete**.

Deleting a schema removes *all* the replication information within that schema, and none of that schema's replications will occur from then on.

The schema node

Expanding a schema node reveals two descriptive nodes that exist only to explain their child nodes. The first node contains the source database, and the second contains the target databases.

Choosing the source database

To specify the database from which replicated data will originate, follow these steps:

1. Expand a schema node and select its Source Database node.
2. Double-click the Add Source Database icon or choose **Replication | Source | Add** to open the Source Database dialog.

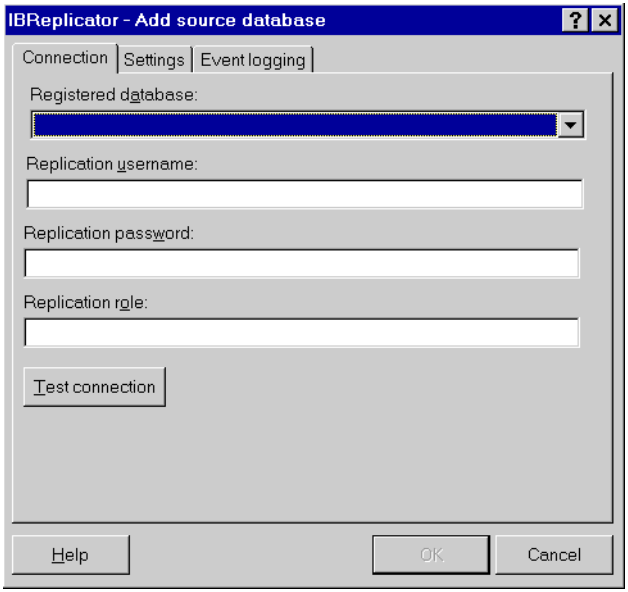

3. Choose one of your registered databases to be the source of the data that is to be replicated (see "Registering databases" on page 12-14).

4 Supply a username and password and/or a rolename for the Replication Server to use when it connects to that database to retrieve the data to be replicated.

   There are also miscellaneous settings for customizing IBReplicator's behavior. These are initialized to default values (see "Customizing default settings" on page 12-16). You can override them for the particular source database that you are defining.

5 Click OK to add a source database node to the Tree view as a child of the Source Databases node.

Selecting a source database node gives access to the commands for editing and removing a source database; these commands can also be accessed by choosing **Replication | Source | Edit** or **Replication | Source | Delete**.

There is also a command for creating system objects. This vitally important command is discussed in "Creating system objects" on page 12-27).

- Each schema can have only one source database, so define a schema for each of your source databases.
- A source database in one schema can be a target database too, but only in another schema.

Choosing the target databases

A target database is one to which replicated data can be sent. When identical data is to be sent to more than one target, it is possible to define one target and then copy or "clone" that definition and apply it to other targets. This section describes these operations.

Specifying a target database

To specify a database to which replicated data can be sent, follow these steps:

1 Expand a schema node and select its Target Databases node.

2 Double-click the Add Target Database icon or choose **Replication | Targets | Add** to display the Target Database dialog.

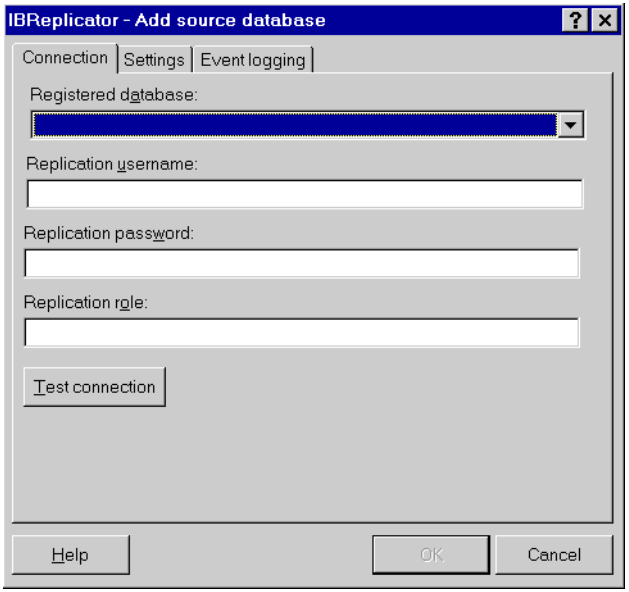

3 Choose one of your registered databases as the source of the data to be replicated (see "Registering databases" on page 12-14).

4 Supply a username and password for the Replication Server to use when connecting to that database to retrieve the data which is to be replicated. You may want to specify a replication role as well. (Compare the *administrative* username, password and role described in "Registering databases" on page 12-14).

   You can also use the Target Database dialog to specify synchronization settings that override the schema's settings for that particular target (see "Creating schemata" on page 12-17).

5 Click OK to add a target database node to the Tree view as a child of the Target Databases node.

Editing and deleting targets Selecting a particular target database node gives access to the commands for editing and removing that target database; these commands can also be accessed by choosing **Replication | Target | Edit** or **Replication | Target | Delete**.

Cloning targets

The tables, columns and rows to be replicated are defined for each target database involved. This allows different data to be replicated to different targets. Typically, however, all the targets in a schema will be receiving similar or identical data, and setting up numerous identical targets can become very tedious indeed.

You can clone targets within a schema. To clone a target, follow these steps:

1 Either drag a target database node onto the schema's Target Databases node, or select a target database in the Tree view and choose **Replication | Target | Clone**.

2 Supply the name of the target database. All information for its replicated tables, columns, and rows is automatically duplicated.

The new target must have the same structure as the target being cloned.

Note A schema can supply one or several targets for its source database, and different data can be replicated to different targets (see "Choosing replicated tables and stored procedures" on page 12-21 and "Choosing replicated columns" on page 12-26 for the details).

Choosing replicated tables and stored procedures

The next step is to identify which tables contain data to be replicated.

Specifying replicated tables

To identify the tables that contain data to be replicated, follow these steps:

1 Expand a node representing a target database and select the Replicated Tables node.

2 Double-click the Replicated Tables icon in Replication Manager's control panel, or choose **Replication | Tables | Define**. This opens the Replicated Tables and Procedures dialog.

3 Map tables in the target database to the tables in the source database by clicking and dragging them.

The above steps can be done automatically by choosing **Auto Generate Tables, Keys and Fields**. In this case, all tables in the Source database are automatically mapped to tables of the same name in the target database. The Manager does not check that the tables exist in the target database, but assumes that source and target databases are identical.

Tables that have already been mapped are ignored by this process.

In cases where source and target databases are not identical, it is often still appropriate to use one of these options:

- Map all tables, and manually remove/modify different tables

or

- First manually map all different tables, then let the automatic process take care of the rest

Removing replicated tables

Use the Replicated Tables dialog to remove source-target table mappings, or choose **Replication | Tables | Remove**.

Note You are free to define complex mappings where source and target tables have different names and even different columns, provided that:

- The two tables both have primary keys that uniquely identify the rows containing data to be replicated.
- The columns containing data to be replicated have compatible datatypes.

Table settings and row-level replication.

To override a source-target table mapping's default settings (see "Customizing default settings" on page 12-16), follow these steps:

1. Select the node in the Tree view that represents the mapping,
2. Double-click the Table Settings icon in Replication Manager's control panel or choose **Replication | Tables | Settings** to open the Table Settings dialog.

Using WHERE conditions to identify rows

You can use the Table Settings dialog to provide a WHERE condition that identifies which rows in the table should be replicated.

The basic syntax of the SQL SELECT statement can be represented informally as:

```
SELECT columnlist
   FROM   table
   WHERE  condition
```

As an example, a request to find all unpaid invoices might look like this:

```
SELECT *
   FROM   SALES
   WHERE  PAID = "n"
```

So, if you want to replicate only rows detailing unpaid invoices, set the Row-level Replication Condition in the Table Settings dialog to:

```
:PAID = "n"
```

This is a very simple condition, but there is no limit to the condition's complexity: any condition that is valid in a SELECT statement is valid as a replication condition.

Important *Test the* condition with a select statement in IBConsole and then copy and paste it into the replication condition's edit box. A syntactically invalid expression is inconvenient rather than disastrous: your log file will be filled with SQL errors. But a valid condition which identifies the wrong rows can indeed be disastrous: consider the case where the condition is being used to replicate public data and leave confidential data behind. Be careful here!

Note Identify each column name in the replication condition with a prefixed colon; this allows IBReplicator to add the `NEW.` and `OLD.` that the system's triggers need to implement row-level replication.

Replicating to stored procedures

Displaying target stored procedures The Replicated Tables dialog shows the tables in the target database by default. To see the target's stored procedures instead, check the Stored Procedures radio button. This allows you to map *target* stored procedures to the *source* tables.

Sophisticated users can define stored procedures to do the work of inserting, updating, or deleting rows on the target database. This allows you to customize IBReplicator to handle specialized cases.

To define a stored procedure for IBReplicator to call, do the following:

- Provide one parameter for each field to be replicated, with data fields listed first, in alphabetical order, and then the key fields, also in alphabetical order. The last parameter defines the action to be taken. It should be a single character, which the server will set when it calls the procedure. The possible values are:

`D`	The row identified by the key fields' values should be deleted
`I`	A row should be inserted with the supplied values in its fields
`U`	The row identified by the key fields' values should be updated to the supplied values

- *Return* an integer value indicating the action's result:

`0`	Success
`1`	• For inserts, a row with the specified key already exists (a primary key violation occurred) • For updates and deletes, no row with the specified key was found
`2`	• For updates and deletes, too many rows were found to act upon • For inserts, an unexpected error that was not a primary key violation occurred.
`-`	• The procedure can also raise an exception for the Replication Server to trap. When an exception is raised, the server rolls back its current transaction and the exception appears in the log, unless it indicates a primary key violation that occurred while inserting; in this case, the conflict resolution rules are used to handle the exception.

In many cases, your procedure will be able to handle both conditions 1 and 2 and will therefore always return zero. For example, an inserted row may be there already, so the procedure can choose to change the insert into an update, or to insert the row somewhere else, or to give it a new primary key. In such cases, the procedure can safely return zero, indicating success.

Example The source database contains this table:

```
CREATE TABLE T (
```

```
  K1 INTEGER NOT NULL,
  K2 VARCHAR(10) NOT NULL,
  K3 DATE NOT NULL,
  F1 VARCHAR(20),
  F2 DOUBLE PRECISION,
  F3 INTEGER,
  F4 VARCHAR(100),
  F5 NUMERIC(4,1),
  F6 NUMERIC(9,2),
  F7 NUMERIC(15,2),
  PRIMARY KEY (K1,K2,K3)
);
```

A stored procedure can then be defined on a target database with the same table as follows:

```
CREATE PROCEDURE REPLICATE_T (
  F1      VARCHAR(20),
  F2      DOUBLE PRECISION,
  F3      INTEGER,
  F4      VARCHAR(100),
  F5      NUMERIC(4, 1),
  F6      NUMERIC(9, 2),
  F7      NUMERIC(15, 2),
  K1      INTEGER,
  K2      VARCHAR(10),
  K3      DATE,
  TYPE    CHAR(1)
) RETURNS (RESULT INTEGER)
AS
  DECLARE VARIABLE COUNTER INTEGER;
BEGIN

  RESULT = 0;  /*default return value*/

  SELECT COUNT(*) FROM T
    WHERE K1 = :K1 AND K2 = :K2 AND K3 = :K3
      INTO :COUNTER;

/* Inserts: If the row already exists, then exit with result=1. This
 * causes the Replication Server to apply conflict rules, which will
 * probably cause the procedure to be called again, but with an
 * with a TYPE of "U". An alternative approach would simply change
 * TYPE to "U" and proceed to update the row instead.
 */
  IF (TYPE = 'I' AND COUNTER > 0) THEN
  BEGIN
    RESULT = 1;
    EXIT;
  END
```

```
/* Updates: If the row does not exist, then exit with result=1. This
 * causes the Replication Server to apply conflict rules, which will
 * probably cause the procedure to be called again, but with a TYPE
 * of "I". An alternative approach would simply change TYPE to "I"
 * and proceed to insert the row instead.
 */
  IF (TYPE = 'U' AND COUNTER = 0) THEN
  BEGIN
    RESULT = 1;
    EXIT;
  END
/* Deletes: If the row does not exist then exit with result=1. The
 * Replication Server will log the error, but otherwise ignore it.
 */
  IF (TYPE = 'D' AND COUNTER = 0) THEN
  BEGIN
    RESULT = 1;
    EXIT;
  END

  IF (TYPE = 'I') THEN
    INSERT INTO T(K1,K2,K3,F1,F2,F3,F4,F5,F6,F7)
    VALUES (:K1,:K2,:K3,:F1,:F2,:F3,:F4,:F5,:F6,:F7);

  IF (TYPE = 'U') THEN
    UPDATE T
    SET    F1 = :F1,
           F2 = :F2,
           F3 = :F3,
           F4 = :F4,
           F5 = :F5,
           F6 = :F6,
           F7 = :F7
    WHERE  K1 = :K1
      AND  K2 = :K2
      AND  K3 = :K3;

  IF (TYPE = 'D') THEN
    DELETE FROM T
    WHERE K1 = :K1
      AND K2 = :K2
      AND K3 = :K3;
EXIT;
END
```

Note InterBase stored procedures do not support Blob and array types as parameters, so such columns cannot be replicated with stored procedures.

Choosing replicated columns

IBReplicator needs to know which of the columns in each replicated table contain the data that should be replicated, and it also needs to know which columns in the tables should be treated as primary keys.

Defining primary keys

1. To identify a replicated table's primary key columns, expand a node representing a source and target table mapping and select the Key Columns node.
2. Double-click the Define Primary Key icon or choose **Replication | Columns | Key** to open the Key Columns dialog
3. To map target and source columns, click and drag the target table's key columns to their corresponding columns in the source table.

Identifying data columns

1. To identify the columns that contain data which is to be replicated, expand one of the nodes representing a source and target table mapping and select the Data Columns node.
2. Double-click the Define Data Columns icon or choose **Replication | Columns | Data** to open the Data Columns dialog.
3. To map columns in the target table to their corresponding columns in the source table, click and drag them across.

Creating system objects

When you have completed a replication schema by specifying a source database and its targets, and mapping the source and target tables and columns to each other, your final step is to tell Replication Manager to create for you the system's tables and triggers in the source database. Select a source database node in the replications Tree view and then double-click the Create System Objects icon in Replication Manager's control panel, or choose **Replication | Source | Create system objects**.

Special user REPL

In a replication source database, any changes made to the database by the user REPL will *not* be replicated. This is useful when you need to make bulk changes to a database, and do not want those changes replicated, as in the following situations:

- Bi-directional replication
- n-way replication

Below are explanations of these special replication situations.

Bi-Directional replication

Bi-Directional replication is used when there are two databases, and each needs to send all changes to the other; they are both source databases, and they are both target databases. The concept is simple, but care needs to be taken to prevent changes from bouncing back and forth in a loop.

In a replication environment, any change to data in a replicated table gets moved from the source database to the target. In a bi-directional setup, this causes problems, because a change made to one database gets sent to the second database, and because it induces a change to the second database, the change gets sent straight back again, in a loop.

Fortunately, there is an easy way to get around this by using the special case user called REPL. This user is not created automatically, but must be created by the DBA, and sufficient privileges must be granted.

Any change made to a database by the user REPL will not get replicated to the target database.

When setting up each *target* database in a bi-directional environment, tell the Replication Server to log in as the user REPL.

n-Way replication

n-Way replication is replication between more than two databases. This can be set up in a number of ways; in its simplest form, each database replicates from/to all others. In this case it is very similar to Bi-Directional Replication and can be set up in exactly the same way.

In more complex scenarios, more work needs to be done. For example, in Hub-and-Spoke setup, a central database controls the replication, fetching data from each database, and forwarding the data to all of the others:

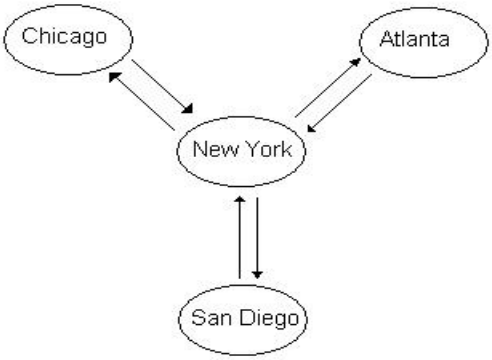

In this example, the New York database fetches data from Chicago, and then forwards it on to Atlanta and San Diego. Clearly, the data from Chicago cannot be written to the New York database using the Special User REPL, otherwise it would stop there, and not be forwarded to Atlanta and San Diego. But, if the user REPL is not used, then the data will bounce straight back to Chicago.

In order to get around this, create a few new usernames and use the 'Row Level Replication Condition' to prevent data from bouncing back to the originating database.

In the case of Chicago to New York, the Replication Server logs in as REPL_CHICAGO. In the reverse case (New York to Chicago), we log in as REPL, and have a Row Level Condition of ' USER <> REPL_CHICAGO'. This means that data will move from Chicago to New York, and from there onwards to Atlanta and San Diego, but will not bounce back to Chicago. Other changes in New York will move to Chicago (amongst others) and stop there because of the use of the user REPL.

Here is a table of the UserNames and conditions required in the above scenario:

Source	Target	UserName	Condition
Chicago	New York	REPL_CHICAGO	
San Diego	New York	REPL_SANDIEGO	
Atlanta	New York	REPL_ATLANTA	

New York	Chicago	REPL	USER < 'REPL_CHICAGO'
New York	San Diego	REPL	USER < 'REPL_SANDIEGO'
New York	Atlanta	REPL	USER < 'REPL_ATLANTA'

This ensures that data is replicated from one branch into the hub in New York. The hub then replicates this data to the other branches, but not back to the originating branch.

Problems and workarounds

The following is a list of problems that some users have encountered along with the techniques for surmounting them.

Problem: Array fields
Array fields are supported, but they are copied across in one unit. This may cause problems for very large arrays.

Workaround
Don't replicate large array fields.

Problem: Computed fields
If you select a computed field for replication, IBReplicator display an error message and does not replicate *any* data for the table in question. This will be fixed in a later release.

Workaround
Don't select computed fields for replication.

Problem: Redundant replication
When only certain of a table's fields have been selected for replication, then data in that table is be replicated, even if the changes affect only fields *not* selected for replication.

Workaround
None needed.

Problem: Floating-point primary keys
Fields that store floating-point numbers should *not* be identified as primary key fields in the replication configuration. The inevitable rounding errors make these fields unsuitable for uniquely identifying the rows in a table.

Workaround
Do not use FLOAT or DOUBLE PRECISION fields as part of the primary key. Ensure that NUMERIC and DECIMAL fields have a precision of at least 9.

Problem: Shared resource error
"Cannot create shared resource" error when trying to run either the replication or the scheduler managers.

Workaround

Shut down the scheduler, and restart it after using the manager.

Appendix

A

Migrating to InterBase 6 and later

InterBase is a mature product that was originally architected before current standards came into existence. As the standards evolved, it became clear that bringing InterBase into compliance with them would produce a somewhat challenging migration path.

With the advent of InterBase 6, Borland decided that it was time to take the leap. InterBase 6 introduced an increased compliance with the SQL-92 standard, but migrating older (InterBase 5 and earlier) clients and databases might, in some cases, require considerable attention to detail.

The feature areas affected are: the use of double quotes, which are now reserved for delimited identifiers; the meaning of the DATE datatype; the behavior of exact numeric datatypes, and the existence of new keywords that might conflict with older metadata names.

This document describes how to plan and execute a smooth migration from earlier versions of InterBase to InterBase 6 or later.

The earlier pages of this guide discuss the issues involved in the migration. Near the end, you will find detailed, step-by-step instructions for both in-place migration and for migrating an old database to a new one. See "**Migrating servers and databases**" on page A-20, "**Migrating databases to dialect 3**" on page A-24, and "**Migrating clients**" on page A-31.

Migration process

These are the steps you must take to migrate servers, databases, and clients. Each is discussed in detail in later sections:

Server and database migration

1. Backup all databases to be migrated
2. Install the latest InterBase server
3. Restore databases to be migrated using the most recent **gbak**; at this point, you have dialect 1 databases
4. Validate migrated databases
5. Migrate databases to SQL dialect 3 (see pages A-24 to A-31)

Client migration

1. Identify the clients that must be upgraded.
2. Identify areas in your application which may need upgrading.
3. Install the InterBase client to each machine that requires it.
4. Upgrade SQL applications to SQL dialect 3.

Migration Issues

Before migrating your databases, you need to learn about InterBase SQL dialects and understand their effect on servers, clients, and the use of certain features introduced in InterBase 6.

InterBase SQL dialects

InterBase recognizes different client and database dialects to allow users more mobility in how their legacy databases are used, accessed, and updated. Beginning with InterBase 6, each client and database has a SQL *dialect*: an indicator that instructs an InterBase 6 or later server how to interpret *transition features*: those features whose meanings have changed between InterBase versions. The following transition features have different meanings based on the dialect used by the client applications:

- Double quote (“): changed from a synonym for the single quote (‘) to the delimiter for an object name.
- DECIMAL and NUMERIC datatypes with precision greater than 9: now stored as INT64 datatypes instead of DOUBLE PRECISION.
- DATE, TIME, and TIMESTAMP datatypes: DATE has changed from a 64-bit quantity containing both date and time information to a 32-bit quantity containing only date information. TIME is a 32-bit quantity containing only time information, while TIMESTAMP is a 64-bit quantity containing both date and time information (the same as DATE in pre-Version 6 SQL).

Clients and databases

Clients and databases each have dialects. Servers do not themselves have a dialect, but they interpret data structures and client requests based on the dialect of each. Applications using an older version of the InterBase client work with InterBase 6 and 7 servers and their databases with some restrictions:

- Version 5 clients cannot access dialect 3 columns that are stored as INT64, TIME, or DATE. (DECIMAL and NUMERIC columns with precision greater than 9 are stored as INT64.)
- Version 5 clients cannot display new datatypes in metadata using the SHOW command, or any equivalent.
- Version 5 clients interpret the DATE datatype as TIMESTAMP, since that was the definition of DATE prior to InterBase 6.
- Version 5 clients cannot access any object named with a delimited identifiers.
- Clients that use the Borland Database Engine (BDE) to access an InterBase 6.0 server are not able to access any of the new field type regardless of the version of the InterBase client installed.

Keywords used as identifiers

Version 5 clients have one advantage over version 6 clients: If you migrate an older database that uses some version 6 keywords as identifiers to version 6 dialect 1, these older version 5 clients can still access those keyword objects. Version 6 dialect 1 cannot do so. Dialect 3 clients can access these keyword objects if the objects are delimited in double quotes.

If version 5 clients use any InterBase 6 or 7 keywords as object names, the InterBase 6 server permits this without error because it recognizes that these clients were created at a time when these were not keywords.

Example For example, the following statement uses the new keyword word TIME:

```
SELECT TIME FROM atable;
```

This statement, when executed via a pre-InterBase 6 client returns the information as it did in previous versions. If this same query is issued using a version 6 or 7 client, an error is returned since TIME is now a reserved word. See page A-9 for a list of new keywords.

Understanding SQL dialects

Below are explanations of server and client behavior with SQL dialects 1, 2, and 3.

Dialect 1 clients and databases

In dialect 1, the InterBase 6 and InterBase 7 servers interpret transition features exactly as an InterBase 5 server does:

- Double quoted text is interpreted as a string literal. Delimited identifiers are not available.
- The DATE datatype contains both time and date information and is interpreted as TIMESTAMP; the name has changed but the meaning has not. Dialect 1 clients expect the entire timestamp to be returned. In dialect 1, DATE and TIMESTAMP are identical.
- The TIME datatype is not available.
- Dialect 1 databases store DECIMAL and NUMERIC datatypes with precision greater than 9 as DOUBLE PRECISION, not INT64.
- Dialect 1 clients expect information stored DECIMAL and NUMERIC datatypes to be returned as double precision; such clients cannot create database fields to hold 64-bit integers.

InterBase 6 and later servers recognize all the other InterBase features in dialect 1 clients and databases.

Dialect 2 clients

Dialect 2 is available only on the client side. It is intended for assessing possible problems in legacy metadata that is being migrated to dialect 3. To determine where the problem spots are when you migrate a database from dialect 1 to dialect 3, you extract the metadata from the database, set **isql** to dialect 2, and then run that metadata file through **isql**. **isql** issues warning whenever it encounters double quotes, DATE datatypes, or large exact numerics to alert you to places where you might need to change the metadata in order to make a successful migration to dialect 3.

To detect problem areas in the metadata of a database that you are migrating, extract the metadata and run it through a dialect 2 client, which will report all instances of transition features. For example:

```
isql -i v5metadata.sql
```

Do not assign dialect 2 to databases.

Dialect 3 clients and databases

In dialect 3, the InterBase server interprets transition features as InterBase 6 SQL 92-compliant:

- Double quoted strings are treated as delimited identifiers.
- Dialect 3 DATE datatype fields contain only date information. Dialect 3 clients expect only date information from a field of datatype DATE.
- The TIME datatype is available, and stores only time information.
- Dialect 3 databases store DECIMAL and NUMERIC datatypes with precision greater than 9 as INT64 *if and only if they are in columns that were created in dialect 3*.

- Dialect 3 clients expect DECIMAL and NUMERIC datatypes with precision greater than 9 to be returned as INT64.

To learn how to migrate older data to INT64 storage, see "Do you really need to migrate your NUMERIC and DECIMAL datatypes?" on page A-28 and "Migrating NUMERIC and DECIMAL datatypes" on page A-29.

Setting SQL dialects

You can set the SQL dialect for a server or client in a variety of ways. For example, the IBConsole user interface has menu options for specifying the SQL dialect. See the *Operations Guide* for a complete explanation of using IBConsole. This section explores the command-line methods for setting a dialect.

Setting the isql client dialect

To use **isql** to create a database in a particular dialect, first set **isql** to the desired dialect and then use it to create the database.You can set **isql** dialect in the following ways:

- On the command line, start **isql** with option **-sql_dialect *n***, where *n* is 1, 2, or 3.

```
isql -sql_dialect n
```

- Within an **isql** session or in a SQL script, you can issue this statement:

```
SET SQL DIALECT n
```

The following table shows the precedence for setting **isql** dialect:

Table A.1 isql dialect precedence

Ranking	How dialect is set
Lowest	Dialect of an attached Version 6 database
Next lowest	Dialect specified on the command line: `isql -sql_dialect n`
Next highest	Dialect specified during the session: `SET SQL DIALECT n;`
Highest	Dialect of an attached Version 5 database (=1)

In InterBase 6, **isql** has the following behavior with respect to dialects:

- If you start **isql** and attach to a database without specifying a dialect, **isql** takes on the dialect of the database.
- If you specify a dialect on the command line when you invoke **isql**, it retains that dialect after connection unless explicitly changed.

- When you change the dialect during a session using SET SQL DIALECT *n*, **isql** continues to operate in that dialect until explicitly changed.
- When you create a database using **isql**, the database is created with the dialect of the **isql** client; for example, if **isql** has been set to dialect 1, when you create a database, it is a dialect 1 database.
- If you create a database without first specifying a dialect for the **isql** client or attaching to a database, **isql** creates the database in dialect 3.

The statements above are true whether you are running **isql** as a command-line utility or are accessing it through IBConsole, InterBase's new interface.

Important Any InterBase 6 **isql** client that attaches to a version 5 database resets to dialect 1.

Setting the gpre dialect

In InterBase 6, **gpre**'s default behavior is to take on the dialect of the database to which it is connected. This enables **gpre** to parse pre-Version 6 source files without moderation.

There are two ways to change the dialect of **gpre**:

- Start **gpre** with option **-sql_dialect *n***. For example, this command sets **gpre** to dialect 3:

```
gpre -sql_dialect 3
```

- Specify dialect within the source, for example:

```
EXEC SQL
   SET SQL DIALECT n
```

The dialect precedence for **gpre** is as follows:

Lowest	Dialect of an attached database
Middle	Command line specification: `gpre -sql_dialect n`

Highest	Dialect explicitly specified within the source, for example
	`EXEC SQL`
	`SET SQL DIALECT n`

Setting the database dialect

To set the dialect of an ODS 10 or later database, attach to the database as either the owner or SYSDBA. Use **gfix** with the command-line option **-sql_dialect *n***, where *n* is 1 or 3. For example, the following statement sets *mydb.gdb* to dialect 3:

```
gfix -sql_dialect 3 mydb.gdb
```

See "Migrating databases to dialect 3" on page A-24 for details about issues to consider before you issue the command.

Features and dialects

Many of the features introduced in InterBase 6 and later operate without reference to dialect. Other features are dialect-specific. The dialect-specific features are discussed below:

Features available in all dialects

The following new features are available in both dialect1 and dialect 3:

IBConsole, InterBase's graphical interface

IBConsole, InterBase's graphical user interface, combines the functionality of the older Server Manager and InterBase Windows ISQL. You now create and maintain databases, configure and maintain servers, and execute interactive SQL from one integrated interface.

Read-only databases

You can make InterBase 6 databases be read-only. This permits distribution on read-only media such as CDROMs and reduces the chance of accidental or malicious changes to databases.

Altering column definitions

The ALTER COLUMN clause of the ALTER TABLE statement can change a column's name, datatype, or position.

Altering domain definitions

ALTER DOMAIN now includes the ability to change the name or datatype of a domain definition.

The EXTRACT() function

The new EXTRACT() function extracts information from the new DATE, TIMESTAMP, and TIME datatypes. In dialect 1, you can use it to extract information from the TIMESTAMP datatype. **Note** "DATE" is new in the sense that it has a different meaning in dialect 3 databases than it did previously.

SQL warnings

The InterBase API function set now returns warnings and informational messages along with error messages in the status vector.

The Services API, Install API, and Licensing API

InterBase now provides three new function libraries. The Services API, which is part of the InterBase client library, provides functions for database maintenance tasks, software activation, requesting information about the configuration of databases and server, and working with user entries in the security database.

New gbak functionality

In InterBase 6, **gbak**'s functionality has been extended. gbak can now perform all of the following actions:

- Back up to multiple files and restore to multiple files
- Perform server-side backups and restores using the **-service** switch
- Set databases to read-only mode when restoring

InterBase Express™ (IBX)

IBX provides native Delphi components for InterBase data access, services, and installation. Borland C++ Builder also can access IBX components.

Features available only in dialect 3 databases

The following features are available only in dialect 3 clients and databases because they conflict with dialect 1 usage.

Delimited identifiers

Identifiers can now be keywords, contain spaces, be case sensitive, and contain non-ASCII characters. Such identifiers must be delimited by double quotes. String constants must be delimited by single quotes.

INT64 data storage

In dialect 3 databases, data stored in DECIMAL and NUMERIC columns is stored as INT64 when the precision is greater than 9. This is true only for columns that are *created* in dialect 3. These same datatypes are stored as DOUBLE PRECISION in dialect 1 and in all older InterBase versions. This change in storage also requires different arithmetic algorithms.

DATE and TIME datatypes

In dialect 3, the DATE datatype holds only date information. This is a change from earlier InterBase versions in which it stored the whole timestamp.

Dialect 3 allows the use of the TIME datatype, which hold only the time portion of the timestamp.

New InterBase keywords

InterBase 6, 6.5, and 7 introduced the following new keywords:

BOOLEAN	HOUR	TIMESTAMP
COLUMN	MINUTE	TRUE
CURRENT_DATE	MONTH	TYPE
CURRENT_TIME	PERCENT	UNKNOWN
CURRENT_TIMESTAMP	ROWS	WEEKDAY
DAY	SECOND	YEAR
EXTRACT	TIES	YEARDAY
FALSE	TIME	

These keywords are reserved words in all dialects.

- Beginning with InterBase 6, you cannot create objects in a dialect 1 database that have any of these keywords as object names (identifiers).
- You can migrate a version 5 database that contains these keywords used as identifiers to version 6 or later dialect 1 without changing the object names: a column could be named “YEAR”, for instance.
  - Version 5 clients can access these keyword identifiers without error.
  - Version 6 clients *cannot* access keywords that are used as identifiers. In a dialect 1 database, you must change the names so that they are not keywords.
  - If you migrate directly to dialect 3, you can retain the names, but you must delimit them with double quotes. To retain accessibility for older clients, put the names in all upper case. Delimited identifiers are case sensitive.
- Although TIME is a reserved word in version 6 dialect 1, you cannot use it as a datatype because such databases guarantee datatype compatibility with version 5 clients.
- In dialect 3 databases and clients, any reserved word can be used as an identifier as long as it is delimited with double quotes.

Delimited identifiers

To increase compliance with the SQL 92 standard, InterBase 6 introduces *delimited identifiers*. An identifier is the name of any database object; for instance a table, a column, or a trigger. A delimited identifier is an identifier that is enclosed in double quotes. Because the quotes delimit the boundaries of the name, the possibilities for object names are greatly expanded from previous versions of InterBase. Object names can now:

- mimic keywords

- include spaces (except trailing spaces)
- be case sensitive

How double quotes have changed

Up to and including version 5, InterBase allowed the use of either single or double quotes around string constants. The concept of delimited identifiers did not exist. Beginning with InterBase 6 (dialect 3), anything in single quotes is interpreted as a string constant and anything in double quotes is interpreted as a delimited identifier. Here is the summary:

- In all versions of InterBase, anything inside single quotes is treated as a string constant.
- In InterBase version 5 and older, anything within double quotes is treated as a string constant, because those versions do not have the concept of a delimited identifier.
- Version 6 dialect 1 is a transition mode that behaves like older versions of InterBase with respect to quote marks: it interprets strings within either single or double quotes as string constants.
- Beginning with version 6 dialect 3, InterBase interprets anything inside double quotes as a delimited identifier. Anything inside single quotes is interpreted as a string constant.
- When InterBase servers version 6 or later detect that the client is dialect 1, they permit client DML (data manipulation) statements to contain double quotes and they correctly handle these as string constants. However, they do not permit double quotes in client DDL (data definition) statements because that metadata would not be allowed in dialect 3. Version 6 servers all insist that string constants be delimited with single quotes when clients create new metadata.

DATE, TIME, and TIMESTAMP datatypes

InterBase 6 dialect 3 replaces the old InterBase DATE datatype, which contains both date and time information, with SQL-92 standard TIMESTAMP, DATE, and TIME datatypes. The primary migration problem exists in the source code of application programs that use the InterBase 5 DATE datatype. In InterBase 6, the DATE keyword represents a date-only datatype, while a Version 5 DATE represents a date-and-time datatype.

Columns and domains that are defined as DATE datatype in InterBase 5 DATE appear as TIMESTAMP columns when the database is restored in InterBase 6. However, a TIMESTAMP datatype has four decimal points of precision, while a Version 5 DATE datatype has only two decimal points of precision.

If you migrate your database to dialect 3 and you require only date or only time information from a TIMESTAMP column, you can use ALTER COLUMN to change the

datatype to DATE or TIME. These columns each take only four bytes, whereas TIMESTAMP and the InterBase 5 DATE columns each take eight bytes. If your TIMESTAMP column holds both date and time information, you cannot change it to an InterBase 6 DATE or TIME column using ALTER COLUMN, because ALTER COLUMN does not permit data loss. Dialect use also enforces certain rules:

- In dialect 1, only TIMESTAMP is available. TIMESTAMP is the equivalent of the DATE datatype in previous versions. When you back up an older database and restore it in version 6, all the DATE columns and domains are automatically restored as TIMESTAMP. DATE and TIMESTAMP datatypes are both available and both mean the same thing in dialect 1.
- In dialect 3, TIMESTAMP functions as in dialect 1, but two additional datatypes are available: DATE and TIME. These datatypes function as their names suggest: DATE holds only date information and TIME holds only time.
- In dialect 3, DATE and TIME columns require only four bytes of storage, while TIMESTAMP columns require eight bytes.

The following example shows the differences between dialect 1 and dialect 3 clients when date information is involved.

Example

```
CREATE TABLE table1 (fld1 DATE, fld2 TIME);
INSERT INTO table1 VALUES (CURRENT_DATE, CURRENT_TIME);
```

Using dialect 1 clients

```
SELECT * FROM table1;

   Statement failed, SQLCODE = -804
   Dynamic SQL Error
   -SQL error code = -804
   -datatype unknown
   -Client SQL dialect 1 does not support reference to TIME datatype

SELECT fld1 FROM table1;

   Statement failed, SQLCODE = -206
   Dynamic SQL Error
   -SQL error code = -206
   -Column unknown
   -FLD1
   -Client SQL dialect 1 does not support reference to DATE datatype
```

Using dialect 3 clients

```
SELECT * FROM table1;

   FLD1          FLD2
   =========== =============
   1999-06-25  11:32:30.0000

SELECT fld1 FROM table1;

   FLD1
```

```
===========
1999-06-25
```

Example

```
CREATE TABLE table1 (fld1 TIMESTAMP);

INSERT INTO table1 (fld1) VALUES (CURRENT_TIMESTAMP);
SELECT * FROM table1;
```

In dialect 1:

```
FLD1
===========
25-JUN-1999
```

In dialect 3:

```
FLD1
==========================
1999-06-25 10:24:35.0000
```

Example

```
SELECT CAST (fld1 AS CHAR(5)) FROM table1;
```

In dialect 1:

```
======
25-JU
```

In dialect 3:

```
Statement failed, SQLCODE = -802
arithmetic exception, numeric overflow, or string truncation
```

Converting TIMESTAMP columns to DATE or TIME

Once you have migrated a database to dialect 3, any columns that previously had the DATE datatype will have the TIMESTAMP datatype. If you want to store that data in a DATE or TIME column, follow these steps:

1. Use ALTER TABLE to create a new column of the desired type.
2. Insert the values from the original column into the new column:

   ```
   UPDATE tablename SET new_field = CAST (old_field AS new_field);
   ```

3. Use ALTER TABLE to drop the original column.
4. Use ALTER TABLE … ALTER COLUMN to rename the new column.

Casting date/time datatypes

InterBase 6 dialect 3 no longer allows the use of the CAST operator to remove the date portion of a timestamp by casting the timestamp value to a character value. When you cast a TIMESTAMP to a CHAR or VARCHAR in dialect 3, the destination type must be at least 24 characters in length or InterBase will report a string overflow exception. This is required by the SQL3 standard.

You can use the CAST() function in SELECT statements to translate between date/time datatypes and various character-based datatypes. The character datatype

must be at least 24 characters in length. You can, however, cast a TIMESTAMP to a DATE and then cast the DATE to a CHAR of less than 24 characters. For example:

```
SELECT CAST (CAST (timestamp_col AS DATE) AS CHAR(10)) FROM table1;
```

It is not possible to cast a date/time datatype to or from BLOB, SMALLINT, INTEGER, FLOAT, DOUBLE PRECISION, NUMERIC, or DECIMAL datatypes.

For more information, refer to "Using CAST() to convert dates and times" in the *Embedded SQL Guide*.

Table A.2 outlines the results of casting *to* date/time datatypes:

Table A.2 Results of casting to date/time datatypes

	To		
Cast From	**TIMESTAMP**	**DATE**	**TIME**
VARCHAR(*n*) CHARACTER(*n*) CSTRING(*n*)	String must be in format YYYY-MM-DD HH:MM:SS.thousands	See below.	String must be in format HH:MM:SS.thousands
TIMESTAMP	Always succeeds	Date portion of timestamp	Time portion of timestamp
DATE	Always succeeds; time portion of timestamp set to 0:0:0.0000	Always succeeds	Error
TIME	Always succeeds; date portion of timestamp set to current date	Error	Always succeeds

Casting DATE to string results in YYYY-MM-DD where "MM" is a two-digit month. If the result does not fit in the string variable a string truncation exception is raised. In earlier versions, this case results in DD-Mon-YYYY HH:mm:SS.hundreds where "Mon" was a 3-letter English month abbreviation. Inability to fit in the string variable resulted in a silent truncation.

Casting a string to a date now permits strings of the form:

```
'yyyy-mm-dd'  'yyyy/mm/dd'  'yyyy mm dd'
'yyyy:mm:dd'  'yyyy.mm.dd'
```

In all of the forms above, you can substitute a month name or 3-letter abbreviation in English for the 2-digit numeric month. However, the order must always be 4-digit year, then month, then day.

In previous versions of InterBase, you could enter date strings in a number of forms, including ones that had only two digits for the year. Those forms are still available in InterBase 6. If you enter a date with only two digits for the year, InterBase uses its "sliding window" algorithm to assign a century to the years.

The following forms were available in earlier versions of InterBase and are still permitted in InterBase 6:

```
'mm-dd-yy' 'mm-dd-yyyy'  'mm/dd/yy'    'mm/dd/yyyy'
'mm dd yy' 'mm dd yyyy'  'mm:dd:yy'    'mm:dd:yyyy'
'dd.mm.yy' 'dd.mm.yyyy'
```

If you write out the month name in English or use a three-character English abbreviation, you can enter either the month or the day first. In the following examples, "xxx" stands for either a whole month name or a three-letter abbreviation. All of the following forms are acceptable:

```
'dd-xxx-yy''dd-xxx-yyyy' 'xxx-dd-yy'   'xxx-dd-yyyy'
'dd xxx yy''dd xxx yyyy' 'xxx dd yy'   'xxx dd yyyy'
'dd:xxx:yy''dd:xxx:yyyy' 'xxx:dd:yy'   'xxx:dd:yyyy'
```

For example, the following INSERT statements all insert the date "January 22, 1943":

```
INSERT INTO t1 VALUES ('1943-01-22');
INSERT INTO t1 VALUES ('01/22/1943');
INSERT INTO t1 VALUES ('22.01.1943');
INSERT INTO t1 VALUES ('jan 22 1943');
```

The following statement would enter the date "January 22, 2043":

```
INSERT INTO t1 VALUES ('01/22/43');
```

Table A.3 outlines the results of casting *from* date/time datatypes:

Table A.3 Results of casting to date/time datatypes

Cast From	To VARCHAR(*n*), CHARACTER (*n*), or CSTRING(*n*)
TIMESTAMP	Succeeds if *n* is 24 or more. Resulting string is in format YYYY-MM-DD HH:MM:SS.thousands.
DATE	Succeeds if *n* is 10 or more. Resulting string is in the format YYYY-MM-DD.
TIME	Succeeds if *n* is 13 or more. Resulting string is the format HH:MM:SS.thousands.

Adding and subtracting datetime datatypes

The following table shows the result of adding and subtracting DATE, TIME, TIMESTAMP, and numeric values. "Numeric value" refers to any value that can be cast as an exact numeric value by the database engine (for example, INTEGER, DECIMAL, or NUMERIC).

Table A.4 Adding and subtracting date/time datatypes

Operand1	Operator	Operand2	Result
DATE	+	DATE	Error
DATE	+	TIME	TIMESTAMP (concatenation)
DATE	+	TIMESTAMP	Error
DATE	+	Numeric value	DATE + number of days: fractional part ignored
TIME	+	DATE	TIMESTAMP (concatenation)
TIME	+	TIME	Error
TIME	+	TIMESTAMP	Error
TIME	+	Numeric value	TIME + number of seconds: 24-hour modulo arithmetic
TIMESTAMP	+	DATE	Error
TIMESTAMP	+	TIME	Error
TIMESTAMP	+	TIMESTAMP	Error
TIMESTAMP	+	Numeric value	TIMESTAMP: DATE + number of days; TIME + fraction of day converted to seconds
DATE	–	DATE	DECIMAL(9,0) representing the number of days
DATE	–	TIME	Error
DATE	–	TIMESTAMP	Error
DATE	–	Numeric value	DATE: number of days; fractional part ignored
TIME	–	DATE	Error
TIME	–	TIME	DECIMAL(9,4) representing number of seconds
TIME	–	TIMESTAMP	Error
TIME	–	Numeric value	TIME: number of seconds; 24-hour modulo arithmetic
TIMESTAMP	–	DATE	Error

Table A.4 Adding and subtracting date/time datatypes

Operand1	Operator	Operand2	Result
TIMESTAMP	–	TIME	Error
TIMESTAMP	–	TIMESTAMP	DECIMAL(18,9) representing days and fraction of day
TIMESTAMP	–	Numeric value	TIMESTAMP: DATE – number of days; TIME: fraction of day converted to seconds

Note Numeric value + DATE, TIME, or TIMESTAMP is symmetric to DATE, TIME, or TIMESTAMP + numeric value.

Using date/time datatypes with aggregate functions

You can use the date/time datatypes with the MIN(), MAX(), COUNT() functions, the DISTINCT argument to those functions, and the GROUP BY argument to the SELECT() function. An attempt to use SUM() or AVG() with date/time datatypes returns an error.

Default clauses

CURRENT_DATE, CURRENT_TIME, and CURRENT_TIMESTAMP can be specified as the default clause for a domain or column definition.

Extracting date and time information

The EXTRACT() function extracts date and time information from databases. In dialect 3, the EXTRACT operator allows you to return different parts of a TIMESTAMP value. The EXTRACT operator makes no distinction between dialects when formatting or returning the information. EXTRACT() has the following syntax:

```
EXTRACT (part FROM value)
```

The value passed to the EXTRACT() expression must be DATE, TIME, or TIMESTAMP. Extracting a part that doesn't exist in a datatype results in an error. For example:

```
EXTRACT (TIME FROM aTime)
```

A statement such as `EXTRACT (YEAR from aTime)` would fail.

The datatype of EXTRACT() expressions depends on the specific part being extracted:

Table A.5 Extracting date and time information

Extract	Resulting datatype	Representing
YEAR	SMALLINT	Year, range 0-5400
MONTH	SMALLINT	Month, range 1-12
DAY	SMALLINT	Day, range 1-31
HOUR	SMALLINT	Hour, range 1-23
MINUTE	SMALLINT	Minute, range 1-59
SECOND	DECIMAL(6,4)	Second, range 0-59.9999
WEEKDAY	SMALLINT	Day of the week, range 0-6 (0 = Sunday, 1 = Monday, and so on)
YEARDAY	SMALLINT	Day of the year, range 1-366

```
SELECT EXTRACT (YEAR FROM timestamp_fld) FROM table_name;
   =======
   1999

SELECT EXTRACT (YEAR FROM timestamp_fld) FROM table_name;
   =======
   1999

SELECT EXTRACT (MONTH FROM timestamp_fld) FROM table_name;
   =======
   6

SELECT EXTRACT (DAY FROM timestamp_fld) FROM table_name;
   =======
   25

SELECT EXTRACT (MINUTE FROM timestamp_fld) FROM table_name;
   =======
   24

SELECT EXTRACT (SECOND FROM timestamp_fld) FROM table_name;
   ============
   35.0000

SELECT EXTRACT (WEEKDAY FROM timestamp_fld) FROM table_name;
   =======
   5

SELECT EXTRACT (YEARDAY FROM timestamp_fld) FROM table_name;
   =======
```

```
  175

SELECT EXTRACT (MONTH FROM timestamp_fld) ||
 '-' || EXTRACT (DAY FROM timestamp_fld) ||
 '-' || EXTRACT (YEAR FROM timestamp_fld) FROM table_name;

  ====================
  6-25-1999
```

DECIMAL and NUMERIC datatypes

The following sections highlight some of the changes introduced by InterBase 6 when dealing with numeric values. They need to be considered carefully when migrating your database from dialect 1 to dialect 3. When considering these issues, keep in mind that in order to make use of the new functionality, the statements must be created with a client dialect setting of 3.

The most notable migration issues involve using the division operator and the AVG() function (which also implies division) with exact numeric operands. *Exact numeric* refers to any of the following datatypes: INTEGER, SMALLINT, DECIMAL, NUMERIC. NUMERIC and DECIMAL datatypes that have a precision greater than 9 are called "large exact numerics" in this discussion. Large exact numerics are stored as DOUBLE PRECISION in dialect 1 and as INT64 in columns created in dialect 3.

Important When you migrate an exact numeric column to dialect 3 it is still stored as DOUBLE PRECISION. The migration does not change the way the data is stored because INT64 cannot store the whole range that DOUBLE PRECISION can store. There is potential data loss, so InterBase does not permit direct conversion. If you decide that you want your data stored as INT64, you must create a new column and copy the data. Only exact numeric columns that are *created* in dialect 3 are stored as INT64. The details of the process are provided in **"Migrating databases to dialect 3"** on page A-24.

You might or might not want to change exact numeric columns to INT64 when you migrate to dialect 3. See **"Do you really need to migrate your NUMERIC and DECIMAL datatypes?"** on page A-28 for a discussion of issues.

Dialect 3 features and changes include

- Support for 64 bit integers.
- Overflow protection. In dialect 1, if the product of two integers was bigger than 31 bits, the product was returned modulo 2^{32}. In dialect 3, the true result is returned as a 64-bit integer. Further, if the product, sum, difference, or quotient of two exact numeric values is bigger than 63 bits, InterBase issues an arithmetic overflow error message and terminates the operation. (Previous versions sometimes returned the least-significant portion of the true result.). The stored procedure **bignum** below demonstrates this.

Operations involving division return an exact numeric if both operands are exact numerics in dialect 3. When the same operation is performed in dialect 1, the result is a DOUBLE PRECISION.

To obtain a DOUBLE PRECISION quotient of two exact numeric operands in dialect 3, explicitly cast one of the operands to DOUBLE PRECISION before performing the division:

```
CREATE TABLE table 1 (n1 INTEGER, n2 INTEGER);

INSERT INTO table 1 (n1, n2) VALUES (2, 3);

SELECT n1 / n2 FROM table1;

   ======================
                        0
```

Similarly, to obtain a double precision value when averaging an exact numeric column, you must cast the argument to double precision before the average is calculated:

```
SELECT AVG(CAST(int_col AS DOUBLE PRECISION))FROM table1;
```

Compiled objects

The behavior of a compiled object such as a stored procedure, trigger, check constraint, or default value depends on the dialect setting of the client at the time the object is compiled. Once compiled and validated by the server the object is stored as part of the database and its behavior is constant regardless of the dialect of the client that calls it.

Example Consider the following procedure:

```
CREATE PROCEDURE exact1 (a INTEGER, b INTEGER) RETURNS (c INTEGER)
AS BEGIN
   c = a / b;
   EXIT;
END;
```

When created by a dialect 1 client:

EXECUTE PROCEDURE *exact 1* returns 1 when executed by either a dialect 1 or dialect 3 client.

When created by a dialect 3 client:

EXECUTE PROCEDURE *exact 1* returns 0 when executed by either a dialect 1 or dialect 3 client.

Example Consider the following procedure:

```
CREATE PROCEDURE bignum (a INTEGER, b INTEGER) RETURNS (c NUMERIC(18,0)
AS BEGIN
   c = a * b;
   EXIT;
```

```
END;
```

When created by a dialect 1 client:

EXECUTE PROCEDURE *bignum* (65535, 65535) returns –131071.0000 when executed by either a dialect 1 or dialect 3 client.

When created by a dialect 3 client:

EXECUTE PROCEDURE *bignum* (65535, 65535) returns `*ERROR* can't access INT64` when executed by a dialect 1 client.

EXECUTE PROCEDURE *bignum* (65535, 65535) returns 4294836225 when executed by a dialect 3 client.

Generators

InterBase 6 generators return a 64-bit value, and only wrap around after 2^{64} invocations (assuming an increment of 1), rather than after 2^{32} as in InterBase 5. Applications should use an *ISC_INT64* variable to hold the value returned by a generator. A client using dialect 1 receives only the least significant 32 bits of the updated generator value, but the entire 64-bit value is incremented by the engine even when returning a 32-bit value to a client that uses dialect 1. If your database was using an INTEGER field for holding generator values, you need to recreate the field so that it can hold 64-bit integer values.

Miscellaneous issues

- IN clauses have a limit of 1500 elements

Resolution If you have more than 1500 elements, place the values in a temporary table and use a SELECT subquery in place of the list elements.

- Arithmetic operations on character fields are no longer permitted in client dialect 3

Resolution Explicitly cast the information before performing arithmetic calculations.

- Using **isql** to select from a TIMESTAMP column displays all information when client dialect is 3.

Resolution In versions of InterBase prior to 6.0, the time portion of a timestamp displayed only if SET TIME ON was in effect. In 6.0 client dialect 3, the time portion of the timestamp always displays.

Migrating servers and databases

You can migrate your servers and applications to InterBase 6 at different times. They are separate migrations. Bear the following issues in mind as you plan your migration:

- Older clients can still access databases that have been migrated to InterBase 6. You must be aware, however, that they cannot access new datatypes or data stored as INT64, and they always handle double quoted material as strings.
- InterBase strongly recommends that you establish a migration testbed to check your migration procedures before migrating production servers and databases. The testbed does not need to be on the same platform as the production clients and servers that you are migrating.

The migration path varies somewhat depending on whether you are replacing an existing server or installing a new server and moving old databases there. Upgrading an existing server costs less in money, but may cost more in time and effort. The server and all the databases you migrate with it are unavailable during the upgrade. If you have hardware available for a new InterBase 6 server, the migration can be done in parallel, without interrupting service more than very briefly. This option also offers an easier return path if problems arise with the migration.

"In-place" server migration

This section describes the recommended steps for replacing an InterBase 5 server with an InterBase 6 server.

1. Shut down each database before backup to ensure that no transactions are in progress.
2. Back up all databases on the version 5 server. Include *isc4.gdb* if you want to preserve your configured user IDs.

   As a precaution, you should validate your databases before backing up and then restore each database to ensure that the backup file is valid.
3. Shut down the version 5 server. If your current server is a Superserver, you are not required to uninstall the server if you intend to install over it, although uninstalling is always good practice. You cannot have multiple versions of InterBase on the same machine. If your current server is Classic, you *must* uninstall before installing InterBase 6.
4. Install the version 6 server.

   Note The install does not overwrite *isc4.gdb* or *isc4.gbk*.
5. Start the new server.
   - On Windows NT, go to Services in the Control Panel and start the InterBase Guardian.
   - On Windows 2000, go to **Control Panel | Administrative Tools | Services** and start the InterBase Guardian.
   - On Windows 98/ME, run the InterBase Guardian application.

- On UNIX/Linux platforms, issue the following command to start the InterBase Superserver as user "interbase":

```
# echo "/usr/interbase/bin/ibmgr -start -forever" | su interbase
```

Note that InterBase can run only as user "root" or user "interbase" on UNIX.

6 To restore the list of valid users, follow these steps:

   a Restore *isc4.gbk* to *isc4_old.gdb*

   b Shut down the server

   c Copy *isc4_old.gdb* over *isc4.gdb*

   d Copy *isc4_old.gbk* over *isc4.gbk*

   e Restart the server

7 Delete each older database file. Restore each database from its backup file. This process creates InterBase 6, ODS 10, dialect 1 databases.

8 Perform a full validation of each database.

After performing these steps, you have an InterBase 6 server and InterBase 6, dialect 1 databases. See "About InterBase 6, dialect 1 databases" on page A-23 to understand more about these databases. See "Migrating databases to dialect 3" on page A-24 for a description of how to migrate databases to dialect 3. See "Migrating clients" on page A-31 for an introduction to client migration.

Migrating to a new server

This section describes the recommended steps for installing InterBase 6 as a new server and then migrating databases from a previous InterBase 5 installation. The process differs only slightly from an in-line install.

1 Back up all databases on the version 5 server. Include *isc4.gdb* if you want to preserve your configured user IDs. Shut down the databases before backup to ensure that no transactions are in progress.

2 Install the version 6 server.

3 Start the new version 6 server.

- On Windows NT/2000, go to Services in the Control Panel and start the InterBase Guardian.
- On Windows 98, run the InterBase Guardian application.
- On UNIX/Linux platforms, issue the following command to start the InterBase Superserver as user "interbase":

```
# echo "/usr/interbase/bin/ibmgr -start -forever" | su interbase
```

Note that InterBase can run only as user "root" or user "interbase" on UNIX.

4 Copy the database backup files to the new server and restore each database from its backup file. This process creates InterBase 6, ODS 10, dialect 1 databases. (Note: In later versions of InterBase, it creates the appropriate current ODS, but always dialect 1.)

   Save your backup files until your migration to dialect 3 is complete.

5 To restore the list of valid users, follow these steps:

   a Restore *isc4.gbk* to *isc4_old.gdb*

   b Shut down the server

   c Copy *isc4_old.gdb* over *isc4.gdb*

   d Copy *isc4_old.gbk* over *isc4.gbk*

   e Restart the server

6 Perform a full validation of each database on the new server.

After performing these steps, you have an InterBase 6 server and InterBase 6, dialect 1 databases. See "About InterBase 6, dialect 1 databases" on page A-23 to understand more about these databases. See "Migrating databases to dialect 3" on page A-24 for a description of how to migrate databases to dialect 3. See "Migrating clients" on page A-31 for an introduction to client migration.

About InterBase 6, dialect 1 databases

When you back up a version 5 database and restore it in InterBase 6, what do you have?

- A version 5 client can access everything in the database with no further changes.
- If there are object names—column or table names, for instance—that include any of the 17 new keywords, you must change these names in order to access these objects with a version 6 dialect 1 client. The new ALTER COLUMN clause of ALTER TABLE makes it easy to implement column name changes.
  - Version 5 clients can still access the columns.
  - Dialect 3 clients can access these columns as long as they delimit them with double quotes.
- The 17 new keywords are reserved words. However, the new datatypes TIME and DATE are not available to use as datatypes. DATE columns have the old meaning—both date and time. The new meaning of DATE—date only—is available only in dialect 3.
- All columns that were previously DATE datatype are now TIMESTAMP datatype. TIMESTAMP contains exactly the information that DATE did in previous versions.
- Exact numeric columns—those that have a DECIMAL or NUMERIC datatype with precision greater than 9—are still stored as DOUBLE PRECISION datatypes. All

arithmetic algorithms that worked before on these columns still work as before. It is not possible to store data as INT64 in dialect 1.

Migrating databases to dialect 3

There are four major areas of concern when migrating a database from dialect 1 to dialect 3:

- Double quotes
- The DATE datatype
- Large exact numerics (for purposes of this discussion, NUMERIC and DECIMAL datatypes that have a precision greater than 9)
- Keywords

The process varies somewhat depending on whether you can create an application to move data from your original database to an empty dialect 3 database. If you do not have access to such a utility, you need to perform an in-place migration of the original database.

Overview

In either method, you begin by extracting the metadata from your database, examining it for problem areas, and fixing the problems.

- If you are performing an in-place migration, you copy corrected SQL statements from the metadata file into a new script file, modify them, and run the script against the original database. Then you set the database to dialect 3. There are some final steps to take in the dialect 3 database to store old data as INT64.
- If you have a utility for moving data from the old database to a newly created empty database, you use the modified metadata file to create a new dialect 3 database and use the utility to transfer data from the old database to the new.

In both cases, you must make changes to the new database to accommodate migrated columns that must be stored as INT64 and column constraints and defaults that originally contained double quotes.

The two methods are described below.

Method one: in-place migration

1. If you have not migrated the database to version 6, dialect 1, do so first. Back up the database again.
2. Extract the metadata from the database using **isql -x**. If you are migrating legacy databases that contain GDML, see "Migrating older databases" on page A-31.

3 Prepare an empty text file to use as a script file. As you fix data structures in the metadata files, you will copy them to this file to create a new script.

Note You could also proceed by removing unchanged SQL statements from the original metadata file, but this is more likely to result in problems from statements that were left in error. Borland recommends creating a new script file that contains only the statements that need to be run against the original database.

For the remaining steps, use a text editor to examine and modify the metadata and script files. Place copied statements into the new script file in the same order they occur in the metadata file to avoid dependency errors.

4 Search for each instance of double quotes in the extracted metadata file. These can occur in triggers, stored procedures, views, domains, table column defaults, and constraints. Change each double quote that delimits a string to a single quote. Make a note of any tables that have column-level constraints or column defaults in double quotes.

Copy each changed statement to your script file, but do not copy ALTER TABLE statements whose only double quotes are in column-level constraints or column defaults.

Important When copying trigger or stored procedure code, be sure to include any associated SET TERM statements.

Quoted quotes If there is any chance that you have single or double quotes *inside* of strings, you must search and replace on a case-by-case basis to avoid inappropriate changes. The handling of quotation marks within strings is as follows:

Table A.6 Handling quotation marks inside of strings

String:	`In "peg" mode`
Double-quoted:	`"In ""peg"" mode"`
Single-quoted:	`'In "peg" mode'`
String:	`O'Reilly`
Double-quoted:	`"O'Reilly"`
Single-quoted:	`'O''Reilly'`

5 In the new script file, search for occurrences of the TIMESTAMP datatype. In most cases, these were DATE datatypes in your pre-6 database. For each one, decide whether you want it to be TIME, TIMESTAMP, or DATE in your dialect 3 database. Change it as needed.

6 Repeat step 5 in the metadata file. Copy each changed statement to your new script file.

7. In the new script file, search for occurrences of reserved words that are used as object names and enclose them in double quotes; that makes them delimited identifiers.
8. Repeat step 7 in the metadata file. Copy each changed statement to your new script file.
9. In each of the two files, search for each instance of a DECIMAL or NUMERIC datatype with a precision greater than 9. Consider whether or not you want data stored in that column or with that domain to be stored as DOUBLE PRECISION or INT64. See "Do you really need to migrate your NUMERIC and DECIMAL datatypes?" on page A-28 for a discussion of issues. For occurrences that should be stored as DOUBLE PRECISION, change the datatype to that. Leave occurrences that you want stored as INT64 alone for now. Copy each changed statement that occurs in the metadata file to your new script file.

Perform the following steps in your new script file:

10. Locate each CREATE TRIGGER and CREATE DOMAIN statement and change it to ALTER TRIGGER or ALTER DOMAIN as appropriate.
11. Locate each CREATE VIEW statement. Precede it by a corresponding DROP statement. For example, if you have a CREATE VIEW *foo* statement, put a DROP VIEW *foo* statement right before it, so that when you run this script against your database, each view first gets dropped and then re-created.
12. If you have any ALTER TABLE statements that you copied because they contain named table-level constraints, modify the statement so that it does nothing except drop the named constraint and then add the constraint back with the single quotes.
13. Check that stored procedure statements are ALTER PROCEDURE statements. This should already be the case.
14. At the beginning of the script, put a CONNECT statement that connects to the original database that you are migrating.
15. Make sure your database is backed up and run your script against the database.
16. Use **gfix** to change the database dialect to 3.

    ```
    gfix -sql_dialect 3 database.gdb
    ```

    Note To run **gfix** against a database, you must attach as either the database owner or SYSDBA.
17. At this point, DECIMAL and NUMERIC columns with a precision greater than 9 are still stored as DOUBLE PRECISION. To store the data as INT64, read "Do you really need to migrate your NUMERIC and DECIMAL datatypes?" on page A-28 and, if necessary, follow the steps in "Migrating NUMERIC and DECIMAL datatypes" on page A-29.
18. Validate the database using either IBConsole or **gfix**.

That's it. You've got a dialect 3 database. There is a little more work to do if you want your NUMERIC and DECIMAL columns with a precision of greater than 9 to be stored as INT64. At this point, they are still stored as DOUBLE PRECISION. To decide whether you want to change they way data in these columns is stored, read "**Do you really need to migrate your NUMERIC and DECIMAL datatypes?**" on page A-28 and "**Migrating NUMERIC and DECIMAL datatypes**" on page A-29.

In addition, there are some optional steps you can take that are described in the following sections, **"Column defaults and column constraints"** and **"Unnamed table constraints"**.

Important If you ever extract metadata from the dialect 3 database that you created using the steps above, and if you plan to use that metadata to create a new database, check to see if the extracted metadata contains double quotes delimiting string constants in column defaults, column constraints, or unnamed table constraints. Change any such occurrences to single quotes before using the metadata to create the new database.

Column defaults and column constraints

The steps above permitted you to retain double quoted string constants in column defaults, column constraints, and unnamed table constraints. This is possible because, once created, InterBase stores them in binary form.

Following the steps above creates a dialect 3 database that is fully functional, but if it contains double quoted string constants in column defaults, column constraints, or unnamed column constraints, inconsistencies are visible when you SHOW metadata or extract it. You can choose to resolve these inconsistencies by following these steps:

1. Back up the database.
2. Examine the metadata to detect each occurrence of a column default or column constraint that uses double quotes.
3. For each affected column, use the ALTER COLUMN clause of the ALTER TABLE statement to give the column a temporary name. If column position is likely to be an issue with any of your clients, change the position as well.
4. Create a new column with the desired datatype, giving it the original column name and position.
5. Use UPDATE to copy the data from old column to the new column:

```
UPDATE table_name
  SET new_col = old_col;
```

6. Drop the old column.

Unnamed table constraints

Read the first two paragraphs under "Column defaults and column constraints" on page A-27 to understand why you don't always need to change constraints with double quotes to single-quoted form, and why you might want to change them.

To bring unnamed table constraints that contain double quotes into compliance with the dialect 3 standard, follow these steps:

1. Back up the database.
2. Examine the metadata to detect each occurrence of an unnamed table constraint that uses double quotes.
3. For each occurrence, use SHOW TABLE to see the name that InterBase has assigned to the constraint.
4. Use ALTER TABLE to drop the old constraint, using the name given in the SHOW TABLE output and add a new constraint. For ease in future handling, give the constraint a name.

   If SHOW TABLE shows that InterBase stores the unnamed constraint as "INTEG_2", then issue the following statement to change the constraint:

```
ALTER TABLE foo
   DROP CONSTRAINT INTEG_2,
   ADD CONSTRAINT new_name
      CHECK (col_name IN ('val1', 'val2', 'val3'));
```

About NUMERIC and DECIMAL datatypes

If you back up a NUMERIC or DECIMAL column with a precision greater than 9 (for example, NUMERIC(12,2)) in an InterBase 5 or earlier database and restore the database as InterBase 6, the column is still stored as DOUBLE PRECISION. Because InterBase does not allow datatype conversions that could potentially result in data loss, you cannot use the ALTER COLUMN statement to change the column datatype from DOUBLE PRECISION to INT64. To migrate a DOUBLE PRECISION column to an INT64 column, you must create a new INT64 column and copy the contents of the older column into it.

In InterBase 6 dialect 3, when you create a NUMERIC or DECIMAL column with a precision of greater than 9, data in it is automatically stored as an INT64 exact numeric.

If you want NUMERIC and DECIMAL datatypes with a precision greater than 9 to be stored as exact numerics, you must take some extra steps after migrating to dialect 3. The following sections tell you how to decide whether you really need to take these steps and how to perform them if you decide you want the exact numerics.

Do you really need to migrate your NUMERIC and DECIMAL datatypes?

As you migrate your databases to dialect 3, consider the following questions about columns defined with NUMERIC and DECIMAL datatypes:

- Is the precision less than 10? If so, there is no issue. You can migrate without taking any action and there will be no change in the database and no effect on clients.
- For NUMERIC and DECIMAL columns with precision greater than 9, is DOUBLE PRECISION an appropriate way to store your data?
  - In many cases, the answer is "yes." If you want to continue to store your data as DOUBLE PRECISION, change the datatype of the column to DOUBLE PRECISION either before or after migrating your database to dialect 3. This doesn't change any functionality in dialect 3, but it brings the declaration into line with the storage mode. In a dialect 3 database, newly-created columns of this type are stored as INT64, but migrated columns are still stored as DOUBLE PRECISION. Changing the declaration avoids confusion.
  - DOUBLE PRECISION may not be appropriate or desirable for financial applications and others that are sensitive to rounding errors. In this case, you need to take steps to migrate your column so that it is stored as INT64 in dialect 3. As you make this decision, remember that INT64 does not store the same range as DOUBLE PRECISION. Check whether you will experience data loss and whether this is acceptable.

Migrating NUMERIC and DECIMAL datatypes

Read "Do you really need to migrate your NUMERIC and DECIMAL datatypes?" on page A-28 to decide whether you have columns in a dialect 1 database that would be best stored as 64-bit integers in a dialect 3 database. If this is the case, follow these steps for each column:

1. Migrate your database to InterBase 6 as described in "Method one: in-place migration" on page A-24.
2. Use the ALTER COLUMN clause of the ALTER DATABASE statement to change the name of each affected column to something different from its original name. If column position is going to be an issue with any of your clients, use ALTER COLUMN to change the positions as well.
3. Create a new column for each one that you are migrating. Use the original column names and if necessary, positions. Declare each one as a DECIMAL or NUMERIC with precision greater than 9.
4. Use UPDATE to copy the data from each old column to its corresponding new column:

```
UPDATE tablename
     SET new_col = old_col;
```

5. Check that your data has been successfully copied to the new columns and drop the old columns.

Method two: migrating to a new database

If you can create a data transfer utility that copies data between databases, the process of migrating a database to dialect 3 is considerably simplified.

Overview Extract the metadata from your database, examine it for problem areas, and fix the problems. Use the modified metadata file to create a new dialect 3 database and use an application to transfer data from the old database to the new.

1. If you have not migrated the database to version 6, dialect 1, do so first. Back up the database again.
2. Extract the metadata from the database using **isql**. If you are migrating a database that contains data structures created with GDML, see "Migrating older databases" on page A-31.

For the following steps, use a text editor to examine and modify the metadata file.

3. Search for each occurrence of the TIMESTAMP datatype. In most cases, these were DATE datatypes in your pre-6 database. Decide whether you want it to be TIME, TIMESTAMP, or DATE in your dialect 3 database. Change it as needed.
4. Find all instances of reserved words that are used as object names and enclose them in double quotes to make them delimited identifiers.
5. Search for each instance of double quotes in the extracted metadata file. These can occur in triggers, stored procedures, views, domains, exceptions, table column defaults, and constraints. Change each double quote to a single quote.
6. Search for each instance of a DECIMAL or NUMERIC datatype with a precision greater than 9. Consider whether or not you want that data stored as DOUBLE PRECISION or INT64. See "Do you really need to migrate your NUMERIC and DECIMAL datatypes?" on page A-28 for a discussion of issues. For occurrences that should be stored as DOUBLE PRECISION, change the datatype to that. Leave occurrences that you want stored as INT64 alone for now.
7. At the beginning of the file, enter SET SQL DIALECT 3. On the next line, uncomment the CREATE DATABASE statement and edit it as necessary to create a new database.
8. Run the metadata file as a script to create a new database.
9. Use your data transfer utility to copy data from the old database to the new dialect 3 database. In the case of a large database, allow significant time for this.
10. Validate the database using **gfix**.
11. At this point, DECIMAL and NUMERIC columns with a precision greater than 9 are still stored as DOUBLE PRECISION. To store the data as INT64, read "Do you really need to migrate your NUMERIC and DECIMAL datatypes?" on page A-28 and, if necessary, follow the steps in "Migrating NUMERIC and DECIMAL datatypes" on page A-29.

Migrating older databases

If you have legacy databases in which some data structures were created with GDML, you may need to extract metadata in a slightly different way.

1. Try extracting metadata as described in Step 2 above and examine it to see if all tables and other DDL structures are present. If they are not, delete the metadata file and extract using the **-a** switch instead of the **-x** switch. This extracts objects created in GDML.
2. You may have to change some of the code to SQL form. For example, the following domain definition

```
CREATE DOMAIN NO_INIT_FLAG AS SMALLINT
   ( no_init_flag = 1 or
   no_init_flag = 0 or
   no_init_flag missing);
```

   needs to be translated to:

```
CREATE DOMAIN NO_INIT_FLAG AS SMALLINT
   CHECK ( VALUE = 1 OR VALUE = 0 OR VALUE IS NULL );
```

3. Some code may be commented out. For example:

```
CREATE TABLE BOILER_PLATE (BOILER_PLATE_NAME NAME,
    DATE DATE,
      CREATED_DATE COMPUTED BY /*  Date */);
```

   needs to be changed to:

```
CREATE TABLE BOILER_PLATE (BOILER_PLATE_NAME NAME,
   "DATE" DATE,
     CREATED_DATE COMPUTED BY "DATE");
```

Migrating clients

To migrate an older client application to InterBase 6, install the InterBase 6 client onto the platform where the client application resides. An InterBase server then recognizes that client as a version 6 dialect 1 client.

It is good practice to recompile and relink the application and make note of field names, datatype use, and so on in the new application. When you recompile, state the dialect explicitly:

```
SET SQL DIALECT n;
```

Important If you have databases that use any of the new version 6 keywords as object identifiers and you are not migrating those databases to dialect 3, you might consider not migrating your version 5 clients. If you migrate them to version 6 dialect 1, you lose the ability to access those keyword columns. See "New InterBase keywords" on page A-9.

When you recompile an existing **gpre** client, you must recompile it with the **gpre -sql_dialect *n*** switch.

There are several paths that permit you to create dialect 3 clients that access all new InterBase 6 features:

- In Delphi 5, make calls to functions in the new InterBase Express (IBX) package. Because the Delphi 5 beta includes InterBase 5, it ships with a version of IBX that does not include calls to the new InterBase 6 Services, Install, and Licensing APIs.
- To write embedded SQL applications that address all InterBase 6 dialect 3 functionality, compile them using **gpre -sql_dialect 3**.

Table A.7 Migrating clients: summary

Client	How to migrate
Older applications such as InterBase version 5 applications	• Dialect is 1; there is no way to change the dialect • A version 5 client application becomes version 6 dialect 1 client whenever the InterBase 6 client is installed on the machine with the client

ISQL	• Issue the command line option:
	`-sql_dialect n`
	• Or issue the command
	`SET SQL DIALECT n;`

GPRE	• Issue the command line option
	`-sql_dialect n`
	• Or issue the command
	`EXEC SQL`
	`SET SQL DIALECT n;`

BDE	All applications use SQL dialect 1. To access InterBase dialect 3 features from Delphi 5, use the IBX components
InterClient	InterBase 6: All applications use SQL dialect 1 InterBase 7 introduced InterClient 3, which is a dialect 3 client
Direct API calls	Set the dialect parameter on *isc_dsql_execute_immediate*(), *isc_dsql_exec_immed2*(), *isc_dsql_prepare*() API calls to the desired dialect value: 1 or 3

IBReplicator migration issues

InterBase 6 contains a new version of IBReplicator that should be used instead of previous versions. It contains new features, described in the *Release Notes* and in the *Operations Guide*, and a few changes which should be addressed when moving from InterBase 5 to interBase 6. If you having been using IBReplicator with previous versions of InterBase, keep these issues in mind:

- If you have any schemas defined that have a source database replicating to more than one target (within the same schema), then you should run the **Create System Objects** command for each of those source databases. In such schemas, more than one entry is placed in the log for any row modified. This does not cause any data errors, but does cause some changes to be replicated more than once.

Note Do not run the **Remove System Objects** command, as this will empty the REPL_LOG table.

- If you have been using older licenses purchased from Synectics Software, those licenses will not work with InterBase 6. You must use the version of IBReplicator for Opensource InterBase, or buy new licenses from Borland Software Corporation for use with the version of IBReplicator for Borland InterBase (the certified commercial version of InterBase).

Migrating data from other DBMS products

If you have a large amount of data in another DBMS such as Paradox, the most efficient way to bring the data into InterBase is to export the data from the original DBMS into InterBase external file format. (See the *Data Definition Guide* for more information about InterBase external files.) Then insert the data from the external files into the internal tables. It is best not to have any constraints on new internal tables; you can validate the database more easily once the data is in InterBase. If constraints do exist, you will need triggers to massage the incoming data.

Appendix

B

InterBase Limits

This appendix defines the limits of a number of InterBase characteristics. The values the following table lists are design limits, and in most cases are further restricted by finite resource restrictions in the operating system or computer hardware.

Various InterBase limits

Table B.1 InterBase specifications

Characteristic	Value
Maximum number of clients connected to one server	There is no single number for the maximum number of clients the InterBase server can serve—it depends on a combination of factors including capability of the operating system, limitations of the hardware, and the demands that each client puts on the server. Assuming a "normal" type of client application that executes database operations from human interaction, and a modern server platform (Pentium 150MHz+, 64MB RAM), expect the InterBase server to comfortably handle up to 150 clients. This is a guideline, not a guarantee. Applications that engage in high levels of contention or that perform complex or high-volume operations could cause the practical number of clients to be fewer. Note also that some operating systems do not have the technology to serve 150 incoming network connections.
Maximum database size	No limit is imposed by InterBase; maximum size is defined by the operating system
Maximum number of files per database	By design, 2^{16} (65,536), because the files are enumerated with an unsigned 16-bit integer. Shadow files count toward this limit. This is a design parameter of InterBase, but most operating systems have a much lower limit on the number of files that a single process can have open simultaneously. In some cases, the OS provides a means to raise this limit. Refer to your OS documentation for the default open files limit, and the means to raise this limit.
Maximum number of cache pages per database	65,536; for the sake of performance, a more practical upper limit would be 10,000. Total size of cache pages should never exceed 50% of memory.
Maximum number of databases open in one transaction	No restriction. The parameters in a transaction parameter buffer comprise a linked list, so there is no limit except that imposed by system resources.
Maximum number of tables per database	By design, 2^{16} (65,536), because tables are enumerated with a 16-bit unsigned integer.
Maximum versions per table	255; then no more metadata changes until the database has been backed up and restored.

Table B.1 InterBase specifications

Characteristic	Value
Maximum row size	64KB. Each Blob and array contributes eight bytes to this limit in the form of their Blob handle. Systems tables (tables maintained by the InterBase engine for system data) have a row size limit of 128KB.
Maximum number of rows and columns per table	By design, 2^{32} rows, because rows are enumerated with a 32-bit unsigned integer per table. Number of columns in a row depends on datatypes used. One row can be 64K. For example, you can define 16,384 columns of type INTEGER (four bytes each) in one table.
Maximum number of indexes per table	By design, 2^{16} (65,536), because indexes per table are enumerated with a 16-bit unsigned integer.
Maximum number of indexes per database	By design, 2^{32}, because you can create 2^{16} tables per database, and each table can have 2^{16} indexes.
Maximum index key size	Starts at 256 bytes for a single-column key, and 200 for multicolumn keys; subtract four bytes for each additional column. Example: a single-column CHAR key can be up to 256 – 4 = 252 bytes; a three-column key must add up to 200 – 12 = 188 bytes. Note that multibyte character sets must fit within the key by counting bytes, not by counting characters. For example, a single-column key using 3-byte UNICODE_FSS characters can have a maximum of (256 – 4) / 3 = 84 characters.
Maximum number of events per stored procedure	No restriction by design, but there is a practical limit, given that there is a limit on the length of code in a stored procedure or trigger (see below).
Maximum stored procedure or trigger code size	48KB of BLR, the bytecode language compiled from stored procedure or trigger language.
Maximum Blob size	The size of the largest single Blob depends on the database page size: 1KB page size: largest Blob is 64MB 2KB page size: largest Blob is 512MB 4KB page size: largest Blob is 4GB 8KB page size: largest Blob is 32GB A Blob is a stream of many segments. The maximum Blob segment size is 64KB.

Table B.1 InterBase specifications

Characteristic	Value
Maximum tables in a JOIN	No restriction by design, but the task of joining tables is exponential in relation to number of tables in the join. The largest practical number of tables in a JOIN is about 16, but experiment with your application and a realistic volume of data to find the most complex join that has acceptable performance.
Maximum levels of nested queries	There is no restriction by design. The practical limit depends on the type of queries you nest. Experiment with your application and a realistic volume of data to find the deepest nested query that has acceptable performance.
Maximum number of columns per one composite index	16
Levels of nesting per stored procedure or trigger	• 750 on Windows platforms • 1000 for UNIX platforms
Maximum size of key in SORT clause	32 KB.
Maximum size of external table file	4 GB on Windows NT/2000/XP; 2 GB on Solaris, Linux, and Windows 98/ME
Range of date values	January 1, 100 a.d. to February 29, 32768 a.d.

Index

A

accessing databases 5-1 to 5-8
activating shadows 6-17
adding
  comments in ISQL script files 10-57
  database file 6-4
  shadow files 6-18
  users 5-9, 5-13 to 5-14
ADMIN_DB 5-3
application development
  InterBase Express 11-23
assigning passwords 5-10, 5-12, 5-13
attachments *See* connections
AUTO mode 6-16
automatic commit of DDL statements 10-12, 10-20
average
  data length 9-11
  fill 9-7

B

backing up databases
  converting external files to internal tables 8-6, 8-14
  metadata only 8-5, 8-14
  options 8-4, 8-14
  preventing sweeping 8-5, 8-14
  upgrading the on-disk structure 8-2
bi-directional replication 12-27
binary
  data 10-28
  files 10-28
blindmeta.sql 5-5
BLOB data
  editing 10-28
  ID numbers, retrieval 10-28
  improving access time 8-10
  saving 10-28
  segment size 11-17
BLOB filters 10-47
BLOBDUMP 10-28
buffered writes vs. forced writes 6-26

C

cache buffers, number of 6-24, 9-9
changing
  character set 10-9
  database page size 8-10, 8-17
  gsec entries 5-14
  user names 5-15
character set, changing 10-9, 10-12
CHECK constraints 10-45, 10-46
checksums
  header page 9-8
  ignoring 8-6, 8-14
client dialect 10-7, 10-23
column headers, changing page length 10-21
commands
  gbak 8-13 to 8-21
  gfix 6-37 to 6-39
  gsec 5-12 to 5-15
  isql 10-25 to 10-26
    displaying 10-29, 10-31
    editing 10-28
    executing 10-8, 10-9, 10-30
comments in ISQL scripts 10-57
COMMIT 10-24, 10-57
commit
  after each table 8-11, 8-17
  automatic 10-12, 10-20
conditional shadows 6-17
configuration parameters 3-13
CONNECT 10-22
connecting to databases 4-7 to 4-9
  denying new connections 6-31
  troubleshooting 4-10 to 4-14
connections
  databases 10-22
  denying 6-31
  monitoring 9-1 to 9-4
  remote servers 4-3, 4-16, 10-21, 10-22
constraints 10-45, 10-46
continuation prompt 10-20
corrupt databases
  backing up by ignoring checksums 8-6, 8-14
  repairing 6-29
CREATE DATABASE 10-20, 10-22
creating
  conditional shadows 6-17
  databases 6-8
  multifile databases 6-4
  shadows 6-14 to 6-17, 8-11
creation date, database 9-10

D

data page slots 9-7
data pages 9-7
database administration 5-9
  adding users 5-13
  changing users 5-10, 5-15

overview 1-6
security utility 5-12 to 5-15
database dialect 10-21
database statistics
gstat 9-12
IBConsole 9-5
databases
accessing 5-1 to 5-8
backing up 8-3, 8-13, 8-14, 8-15
closing 4-8, 10-29, 10-32
connecting to 4-7 to 4-9
creating 6-8
creating multifile 6-4
creation date 9-10
deactivating indexes 8-12
deleting 6-10
denying new connections 6-31
disabling validity checking 8-12, 8-17
dropping 6-10
extracting metadata 8-5, 8-14
file naming conventions 6-3
file, adding 6-4
gbak and gsplit 1-10
locking statistics 9-15
name information 9-8
naming 5-3, 6-3
overwriting 8-11, 8-17
page size 8-10 to 8-11, 8-17, 9-8
read-only 6-7, 6-38, 8-17
registering 4-5
repairing 6-29 to 6-30
replacing 8-11, 8-17
restarting 6-30, 6-33
security 5-8
shadowing 6-12 to 6-18, 8-11
shutdown and restart 6-30
shutting down 6-31
structure, retrieving 10-45
sweeping 6-21 to 6-23, 6-38
immediately 6-23, 6-38
sweep interval 6-22
testing connections 4-16
unregistering 4-9
upgrading the on-disk structure 8-2
validating 6-26 to 6-29, 6-39
viewing information 10-45
viewing metadata 10-15
DDL
automatic commits 10-12, 10-20
defined 10-25
extracting 10-20, 10-21, 10-24 to 10-25
processing 10-19
transactions 10-24
deactivating indexes 8-12, 8-17
deadlocks 9-16
DECLARE FILTER 10-18, 10-24, 10-48
default terminator 10-20, 10-33
deleting
databases 6-10
shadows 6-17
users 5-11
dialect 10-23
database 10-21
disabling
garbage collection 8-5, 8-14
validity checking 8-12, 8-17
disabling automatic internet dialup 4-13
displaying
error messages 10-31
metadata 10-15, 10-26
objects 10-51
query plan 10-12, 10-40
statements 10-20
statistics 10-41
version numbers 10-21
DML 10-25
processing 10-19
domains 10-46
DROP SHADOW 6-17
dropping
databases 6-10
shadows 6-17
users 5-11

E

echoing 10-20
EDIT 10-28
INPUT and 10-30
editing
BLOB data 10-28
input files 10-45
isql commands 10-28
editors 10-29
environments, isql 10-26
errors
connection refused 4-10, 4-14
connection rejected 4-12
gbak 8-22, 8-26
gfix 6-39
gsec 5-16
isql 10-26
exceptions 10-47
EXIT 10-29
QUIT vs. 10-32
external tables 6-2
EXTERNAL_FILE_DIRECTORY 6-2
extracting

DDL 10-20, 10-21
metadata 8-5, 10-18, 10-24 to 10-25

F

files
input 10-9, 10-20
naming 5-3, 6-3
naming conventions 6-3
shadow 6-14, 6-18, 8-11
writing to 10-10, 10-21, 10-24
fill distribution 9-7, 9-11
filters, displaying 10-16, 10-47
forced writes vs. buffered writes 6-26

G

garbage collection, disabling 8-5, 8-14
gbak
commands 8-13 to 8-21
errors 8-22 to 8-26
generators, displaying 10-16, 10-49
gfix
activating shadows 6-17
commands 6-37 to 6-39
errors 6-39
killing shadows 6-17
global header information 9-16
gpre
licensing 7-6
gsec 5-12 to 5-15
adding users 5-13 to 5-14
changing entries 5-14
commands 5-12 to 5-15
deleting entries 5-15
errors 5-16
exiting 5-12
help 5-13
options 5-13
running remotely 5-12
starting 5-12, 5-15
gsplit 1-10
gstat 9-12
Guardian 3-1

H

header page generation 9-8
help
gsec 5-13
IBConsole 2-3
InterBase 2-3
ISQL commands 10-29
UNIX 3-6

I

ibconfig 3-12, 3-13
IBConsole
character sets 10-12
commit and rollback 10-10
denying new connections 6-31
displaying metadata 10-15
executing SQL statements 10-8 to 10-11
extracting metadata 10-18
restarting databases 6-33
saving ISQL input and output 10-10
script file, adding comments 10-57
security 5-9
toolbar 2-4
Tree pane 2-5
viewing statistics 9-5
Work pane 2-6
ignoring
checksums 8-6, 8-14
limbo transactions 8-6, 8-14
implementation ID 9-9
index root page 9-7
indexes
correcting duplicate values 8-12
deactivating 8-12, 8-17
depth 9-10
displaying 10-26, 10-50
improving performance 8-10 to 8-11, 11-18
retrieving 10-50
INPUT 10-30
input files 10-9, 10-20, 10-30
editing 10-45
interactive SQL. See isql
InterBase
API, user administration 5-11
developing database applications with IBX 11-23
Guardian 3-1
on UNIX 3-6
version numbers 10-21, 10-54
ISC4.GDB 5-12
viewing contents 5-13, 5-15, 8-20
isql 10-19 to 10-27
commands 10-25 to 10-26
connecting to databases 10-22
displaying help 10-29
editing 10-28
errors 10-26
executing 10-8, 10-9, 10-30
exiting 10-22, 10-26, 10-29, 10-32
options 10-19 to 10-22
output 10-31
saving input and output 10-10

script files, adding comments in IBConsole 10-57
SET statements 10-32 to 10-44
setting environment 10-26
setting SQL dialect 10-23
specifying database 10-19
starting 10-19 to 10-20, 10-22
terminator characters 10-20, 10-21

K

killing shadows 6-17

L

leaf buckets 9-10
limbo transactions
ignoring 8-6, 8-14
two-phase commit 6-33
local access 1-9
lock
manager 1-9, 9-15
table 9-16
locks 9-15
logging in to a server 4-3

M

MANUAL mode 6-16
max dup 9-11
metadata
command-line ISQL 10-19
displaying 10-15, 10-26
extracting 8-5, 10-18, 10-24 to 10-25
in IBConsole 10-15
monitoring attachments 9-1 to 9-4
multifile databases 6-4
multiprocessor support 3-3

N

name information 9-8
naming
databases 5-3, 6-3
nesting INPUT commands 10-30
next
connection ID 9-9
header page 9-10
transaction 9-9
node names, shared disks vs. 10-23
nodes 9-10
number of cache buffers 9-9
n-Way replication 12-28

O

objects
deleting 10-51
displaying 10-51
ODS. See on-disk structure
oldest
active transaction 9-9
transaction 9-9
on-disk structure
upgrading 8-2
version 9-8
operating system shells 10-44
OUTPUT 10-31
output
files 10-10, 10-21, 10-24
isql 10-10, 10-31
metadata 10-15, 10-26
redirecting 10-10, 10-31
statements 10-20
user-defined functions 10-16, 10-48
verbose 8-7, 8-13, 8-15

P

page size
changing 8-10 to 8-11, 8-17
default 6-10
displaying current 9-8
passwords
assigning 5-10, 5-12, 5-13
connecting to remote servers 10-21
PLAN 11-22
PLAN option 10-12, 10-41
primary pointer page 9-7
procedures, listing 10-18, 10-26, 10-51

Q

query 10-19
displaying plan 10-12, 10-40
testing 11-26
QUIT 10-32
EXIT vs. 10-29

R

readmeta.sql 5-5
read-only databases 6-7, 6-38, 8-17
registering
databases 4-5
servers 4-2
remote servers, connecting to 4-3, 4-16, 10-21, 10-22
repairing databases 6-29 to 6-30
replacing databases 8-11, 8-17

replication
bi-directional 12-27
n-Way 12-28
restarting databases 6-30, 6-33
restore options 8-10
ROLLBACK 10-24

S

saving ISQL input and output 10-10
security
adding a user 5-9, 5-13
database 5-8
displaying privileges 10-49
dropping users 5-11, 5-15
IBConsole 5-9
modifying user configuration 5-10, 5-14
security database 5-1
name 5-3
sequence number 9-9
servers
log 3-17
logging in 4-3
registering 4-2
starting on UNIX 3-6
unregistering 4-5
SET 10-32
SET AUTODDL 10-34
SET BLOBDISPLAY 10-35
SET COUNT 10-37
SET ECHO 10-38
SET LIST 10-39
SET NAMES 10-40
SET PLAN 10-40 to 10-41
SET statements 10-32 to 10-44
SET STATS 10-41
SET TERM 10-42 to 10-43
SET TERMINATOR 10-20
SET TIME 10-43
SET TRANSACTION 10-24
shadow count 9-9
shadows
activating 6-17
adding files 6-18
advantages 6-13
AUTO mode 6-16
conditional 6-17
creating 6-14 to 6-17, 8-11
dropping 6-17
killing 6-17
limitations 6-13
MANUAL mode 6-16
overview 6-12 to 6-14
shared disks, node names vs. 10-23
SHELL 10-44
SHOW CHECK 10-45
SHOW DATABASE 6-15, 10-45
SHOW DOMAINS 10-46
SHOW EXCEPTIONS 10-47
SHOW FILTERS 10-47
SHOW FUNCTIONS 10-48
SHOW GENERATORS 10-49
SHOW GRANT 10-49
SHOW INDEX 10-50
SHOW PROCEDURES 10-51
SHOW SYSTEM 10-52
SHOW TABLES 10-53
SHOW TRIGGERS 10-53
SHOW VERSION 10-54
SHOW VIEWS 10-55
shutting down databases
denying new connections 6-31
denying new transactions 6-32
forced shutdown after timeout 6-33, 6-38
timeout options 6-31
SMP support 3-3
source database 12-5
SQL dialect 6-10, 10-21, 10-23
SQL statements
committing in IBConsole 10-10
executing in IBConsole 10-8 to 10-11
rolling back in IBConsole 10-10
SQLCODE 10-26
starting
gsec 5-12, 5-15
InterBase Guardian 3-5
InterBase Server 3-5
isql 10-19 to 10-20, 10-22
statements
displaying 10-20
terminator characters 10-20, 10-21
statistics
displaying 10-41
gstat 9-12
IBConsole 9-5
SuperServer architecture 1-8
sweeping databases 6-21 to 6-23, 6-38
disabling 6-23
immediately 6-23, 6-38
preventing during a backup 8-5, 8-14
sweep interval 6-22
synchronous replication 12-5
SYSDBA 5-2
system
editors 10-29
shells 10-44
tables 10-52
views 10-52

system table security 5-5
system temporary tables 9-1 to 9-4

T

tables
  constraints 10-45, 10-46
  listing 10-16, 10-26
target database 12-5
TCP/IP 10-22
terminator characters
  default 10-20
  isql 10-20, 10-21
testing
  database connection 4-16
text
  editors 10-29
  saving blocks to file 10-28
toolbar 2-4
total dup 9-11
transactions
  committing 10-24, 10-29
  DDL 10-24
  denying new transactions 6-32
  isql 10-24
  oldest 9-9
  oldest active 9-9
  rolling back 10-24, 10-32
  rolling back limbo transactions 8-6
  two-phase commit 6-33
Tree pane 2-5
triggers, listing 10-26, 10-53
troubleshooting, database connection 4-10 to 4-14
trusted host 5-2
two-phase commit 6-33, 8-6

U

unregistering
  databases 4-9
  servers 4-5
upgrading the on-disk structure (ODS) 8-2
use all space 8-12
user administration with the InterBase API 5-11
user names 10-21
  adding 5-12, 5-13 to 5-14
  changing 5-15
  dropping 5-11
  UNIX 5-2
user-defined functions
  listing 10-16, 10-48
  viewing 10-48

V

validating databases 6-26 to 6-29, 6-39
validity checking, disabling 8-12
verbose output 8-7, 8-13, 8-15
version
  numbers, displaying 10-21, 10-54
  on-disk structure 9-8
views, listing 10-16, 10-55

W

warnings 10-21
Windows NT commands 10-44
Work pane 2-6
writemeta.sql 5-5