DBMaker

Database Administrator's Guide

CASEMaker®

CASEMaker Inc./Corporate Headquarters

1680 Civic Center Drive

Santa Clara, CA 95050, U.S.A.

www.casemaker.com

www.casemaker.com/support

Document No. 645049-231178/DBM42-M07232004-DBAG

Publication Date: 2004-07-23

For a description of updated functions that do not appear in this manual, read the file named README.TXT after installing the CASEMaker DBMaker software.

Trademarks

Notices

Contents

1 Introduction

Welcome to the DBMaker Database Administrator's Guide. DBMaker is a powerful and flexible SQL Database Management System (DBMS) that supports an interactive Structured Query Language (SQL), a Microsoft Open Database Connectivity (ODBC) compatible interface, and Embedded SQL for C (*E*SQL/C). DBMaker also supports a Java Database Connectivity (JDBC) compliant interface and DBMaker COBOL interface (DCI). The unique open architecture and native ODBC interface give the user freedom to build custom applications using a wide variety of programming tools, or query a database using existing ODBC, JDBC, or DCI compliant applications.

DBMaker is easily scalable from personal single-user databases to distributed enterprise-wide databases. Regardless of the configuration chosen for a database, the advanced security, integrity, and reliability features of DBMaker ensure the safety of critical data. Extensive cross-platform support helps leverage existing hardware, and allows for expansion and upgrading to more powerful hardware as needs grow.

DBMaker provides excellent multimedia-handling capabilities, allowing all types of multimedia data to be stored, searched, retrieved, and manipulated. Binary Large Objects (BLOBs) ensure the integrity of multimedia data by taking full advantage of the advanced security and crash recovery mechanisms included in DBMaker. File Objects (FOs) allow multimedia data to be managed while maintaining the capability to edit individual files in a source application.

This book is intended for database administrators who are not familiar with either the concepts or principles of the DBMaker DBMS or the syntax and grammar of the

DBMaker SQL query language. However, you should have a general working knowledge of computers, and should be comfortable using the operating system you are using to run DBMaker. Information on the operating system is beyond the scope of this manual; consult your operating system documentation if any problems in this area are encountered.

This book contains general information on the concepts and principles a database administrator should understand when using the DBMaker DBMS. It gives an overview of how to use the DBMaker SQL commands necessary to create, maintain, and optimize a database. To present the information more clearly, examples and illustrations are provided throughout the manual.

The implementation of a DBMS can greatly affect the performance of database operations. It requires many decisions to optimize and tune database performance, such as where to store data, how to access data, whether to have an index, and how to protect the data. This manual provides background to enable clear understanding of the effects of the choices database administrators or application developers make. SQL commands will be used to illustrate most of the functions that DBMaker supports, and references to other database administration tools are provided.

Most of the concepts, commands, and examples in this book are presented using dmSQL, the command-line tool provided with DBMaker. In a few cases some database administration functions can only be performed using one of the other DBMaker application tools or utilities. For more information on how to use the application tools and utilities provided with DBMaker, refer to section 1.1 *Other Sources of Information*.

1.1 Other Sources of Information

DBMaker provides many other user's guides and reference manuals in addition to this reference.

For more information on a particular subject, consult one of these books:

- For more information on the SQL language implemented by DBMaker, refer to the "*SQL Command and Function Reference*".
- For more information on the ESQL/C language implemented by DBMaker, refer to the "*ESQL/C Programmer's Guide*".
- For more information on using dmSQL, refer to the "*dmSQL User's Guide*".
- For more information on error and warning messages, refer to the "*Error and Message Reference*".
- For more information on configuring and managing databases using DBMaker's J Tools, refer to the "*JDBA Tool User's Guide*", "*JServer Manager User's Guide*", and "*JConfiguration Tool Reference*".

1.2 Technical Support

CASEMaker provides thirty days of complimentary email and phone support during the evaluation period. When software is registered an additional thirty days of support will be included. Thus, extending the total support period for software to sixty days. However, CASEMaker will continue to provide email support for any bugs reported after the complimentary support or registered support has expired (free of charges).

Additional support is available beyond the sixty days for most products and may be purchased for twenty percent of the retail price of the product. Please contact sales@casemaker.com for more details and prices.

CASEMaker support contact information for your area (by snail mail, phone, or email) can be located at: www.casemaker.com/support. It is recommended that the current database of FAQ's be searched before contacting CASEMaker support staff.

Please have the following information available when phoning support for a troubleshooting enquiry or include the information with a snail mail or email enquiry:

- Product name and version number
- Registration number
- Registered customer name and address
- Supplier/distributor where product was purchased
- Platform and computer system configuration
- Specific action(s) performed before error(s) occurred
- Error message and number, if any
- Any additional information deemed pertinent

1.3 Document Conventions

This book uses a standard set of typographical conventions for clarity and ease of use. The NOTE, Procedure, Example, and CommandLine conventions also have a second setting used with indentation.

CONVENTION	DESCRIPTION
Italics	Italics indicate placeholders for information that must be supplied, such as user and table names. The word in italics should not be typed, but is replaced by the actual name. Italics also introduce new words, and are occasionally used for emphasis in text.
Boldface	Boldface indicates filenames, database names, table names, column names, user names, and other database schema objects. It is also used to emphasize menu commands in procedural steps.
KEYWORDS	All keywords used by the SQL language appear in uppercase when used in normal paragraph text.
SMALL CAPS	Small capital letters indicate keys on the keyboard. A plus sign (+) between two key names indicates to hold down the first key while pressing the second. A comma (,) between two key names indicates to release the first key before pressing the second key.
NOTE	Contains important information.
➲ **Procedure**	Indicates that procedural steps or sequential items will follow. Many tasks are described using this format to provide a logical sequence of steps for the user to follow
➲ **Example**	Examples are given to clarify descriptions, and commonly include text, as it will appear on the screen. Other forms of this convention include Prototype and Syntax.
`CommandLine`	Indicates text, as it should appear on a text-delimited screen. This format is commonly used to show input and output for dmSQL commands or the content in the dmconfig.ini file

Table 1-1 Document Conventions

2 Overview

The physical organization of data across the files that make up a database can become quite complex. A DBMS, such as DBMaker, isolates a view of the data from the implementation of the database on a computer. The database is viewed as a collection of two-dimensional tables that contain data values in rows and columns. These tables are easy to visualize and flexible for data modeling.

DBMaker provides a number of ways to retrieve the data from tables. The interactive dmSQL line command tool is useful for daily transaction processing or ad-hoc queries, and the DBMaker *application program interface (API)* is ideally suited for developing applications quickly and easily. DBMaker also comes with easy-to-use graphical tools that are consistent across platforms.

2.1 Features

As an SQL database management system, DBMaker has all of the features traditionally found in a relational database management system. DBMaker is also enhanced with many powerful and advanced features. These enhanced features not only increase performance, but also provide DBMaker with capabilities not normally found in traditional database management systems, especially in the area of multimedia support.

Multimedia Support

Powerful multimedia management capabilities built into the database engine allow for efficient storage and manipulation of large amounts of multimedia data including text, graphics, audio, video, and animations. These multimedia management capabilities also provide a great deal of flexibility, allowing multimedia data to be stored in different ways depending on the needs of the user.

Multimedia features include:

- Binary Large OBjects (BLOBs) and File Objects (FOs).
- Multiple BLOB and FO columns in a table.
- File Objects can be edited with existing multimedia tools.
- Built-in full-text search engine.

Multimedia data can be stored directly in the database as Binary Large Objects (BLOBs). This data is protected by the full spectrum of security, reliability, and integrity features provided for conventional data types. In addition, multimedia data can be stored as file objects, allowing third-party multimedia tools full access to multimedia data while keeping it under database control.

JDBC Support

DBMaker supports features of JDBC 3.0 as well as Java Transaction API (JTA) functions. JDBC JTA allows for connections to popular Java AP servers such as BEA WebLogic™.

To learn about implementing JDBC and JDBA, refer to the product documentation. Information about the JDBC specification is available at: http://java.sun.com/products/jdbc/.

Information about the JTA specification is available at; http://java.sun.com/products/jta/.

Microsoft Transaction Server (MTS) Support

MTS (Microsoft Transaction Server) is an integral part of Windows NT, and is installed by default as part of the operating system in Windows2000. MTS evolved as a Transaction Processing system to provide on Windows NT the same kinds of features available on other platforms like CICS, Tuxedo, etc. These are purely designed for creating stable environments for data sources.

DBMaker supports MTS so users can perform transactional operations.

The following are required to use DBMaker with MTS:

- DBMaker requires Microsoft Data Access Components (MDAC) version 2.6 or higher to run with MTS. The latest version of MDAC can be downloaded from http://www.microsoft.com/data.
- If using MDAC 2.5, add the **DM_DifEn = 0** option in the **DM_COMMON_OPTION** section of the **dmconfig.ini** file.
- In the **dmconfig.ini** file, set **DB_DifCo = 1** (default setting) in the database sections that will run with MTS.

Open Interface

Using the native ODBC 3.0 compatible interface and ANSI SQL-99 support, high-performance applications can quickly be created. Applications can be built using a wide variety of popular development tools, including Visual C++, Visual Basic, Delphi, and AcuBench. DBMaker allows developers and administrators to use the tools they already have, and does not restrict them to a proprietary development environment.

Open interface features include:

- ANSI-92 entry-level compliance
- ODBC 3.0 support
- ESQL/C preprocessor

The included ESQL/C preprocessor simplifies the development process for programs written using a traditional C development environment. A database application can be written using the power of the high-level Embedded SQL query language, and the DBMaker preprocessor will automatically translate it to the appropriate ODBC function calls.

Data Integrity

DBMaker also provides a full range of traditional data integrity features. Primary and foreign keys ensure data integrity, with full support for referential actions. User-defined data types, together with domain, column, and table constraints ensure only valid values can be entered in each field.

Data integrity features include:

- Integrity checking of primary and foreign keys
- Full support for referential actions
- Table and column constraints
- User-defined data types
- Default column values

Data Reliability

Data will always be safe thanks to advanced data protection facilities, such as automatic crash recovery, database consistency checking, and automatic backups. These features ensure data consistency and safety in the event of operating system or disk failures.

Data reliability features include:

- Online transaction processing
- Online full and incremental backup
- Automatic crash recovery

- Automatic incremental backup
- Automatic statistic updates
- Database consistency checking
- Multiple journal files
- Optional BLOB backup

Storage Management

Modern storage management facilities provide flexible data storage with simple management and configuration. There is no practical limit on the number of rows in a table, or on the number of tables in a database. A table can even be spread over multiple disks! DBMaker also allows table schema to be altered online, resulting in the ability to develop applications that can dynamically adjust to user needs.

Storage management features include:

- Autoextend and regular tablespaces
- Raw device support on UNIX
- Maximum database size of 256PB (petabytes; 1 PB = 1024TB)
- No practical limit on the number of tables in a database
- No practical limit on the number of records in a table
- Online table schema redefinition

DBMaker can dynamically extend the storage space of a database, up to the limits of available disk space. Storage space may also be fixed and manually adjusted. On UNIX platforms, DBMaker supports raw devices, which allows the file system to be bypassed and data to be written directly to the raw device for maximum performance.

Security Management

The centralized and multi-user nature of a DBMS requires that some form of security control be in place, both to prevent unauthorized access and to limit access for authorized users. User- and group-level authority levels control that accesses a database. Privilege management on tables or individual columns controls what they access.

Security management features include:

- User- and group-level security
- Nested groups
- Privilege management on both tables and individual columns
- Privilege management on stored commands and stored procedures
- Network encryption

Advanced Language Features

Advanced language features complement traditional database functions. Easily extend and customize the capabilities of DBMaker using stored commands, stored procedures, Triggers, and user-defined functions. Business logic can be written directly into the database engine, centralizing the logic in the database so it is easier to manage and maintain.

Advanced language features include:

- Built-in functions
- User-defined functions
- Stored commands
- Stored procedures
- Triggers

2.2 Database Modes

The database administrator may start a database in one of several different database modes. Each mode provides different options for connecting to and accessing a database. This gives the ability to scale a database from a simple single-user system on one computer to a large multi-user system distributed across several computers.

The database modes available depend on the platform a database server runs on, and the connection. DBMaker has three different database modes: single-user, multiple-connection, and client/server.

Single-User Mode

Single-user mode is only available on the UNIX/Linux platforms. This is a simplified version of DBMaker for non-sharable databases. The main advantages of this mode are the smaller application size and faster execution speed for most database operations, since locks, security, and network support are not required for a single user database.

A limitation of this mode is that since only one connection can be made to the database at a time, the database cannot run any of the extra servers or daemons, such as backup server, replication server, or global transaction server. Another limitation is that the database must be accessed from the host machine, since the database is not available over the network.

Multiple-Connection Mode

Multiple-connection mode is only available on the Windows platform. One advantage of this mode is that multiple connections to a database are available, with the full range of security and reliability features of DBMaker. All connections must access the database from the host machine, since there is no network support.

A limitation of this mode is that the database does not support any of the extra servers or daemons, such as backup server, replication server, or global transaction server.

Client/Server Mode

Client/server mode is available on all platforms. This mode permits multiple connections to a database from any computer connected to the host computer via a TCP/IP network, and provides the full range of security, reliability, and concurrency control features of DBMaker. In addition, data sent across the network can be encrypted for additional security. This mode supports all of the extra servers and daemons, such as backup server, replication server, and global transaction server.

2.3 DBMaker Interface and Tools

DBMaker comes complete with an application program interface and several tools and utilities that can be used to manage a database. These tools range from a command-line based interactive SQL query tool, to a graphical tool for managing multiple servers. Novice database users will appreciate the simple management features and graphical tools that are consistent across platforms.

Application Program Interface

The application program interface (API) is a library of low-level routines that operate directly on the database engine. The API is used when creating software applications with a general-purpose programming language such as C++ or Visual Basic. DBMaker provides an ODBC 3.0 compatible interface, and currently supports all core-level functions and most extended-level functions.

dmSQL Interactive Query Tool

dmSQL is a character-based, interactive interface that utilizes the full power and functionality of DBMaker directly. Use dmSQL to manipulate a database, perform ad-hoc queries, and see the result sets immediately. dmSQL is often the only method of exploiting the full power of a database without having to create programs using a conventional programming language.

JDBA Tool

JDBA Tool is a graphical, interactive tool used to maintain and monitor a database. JDBA Tool hides the complexity of the DBMS and query language behind an intuitive, easy to understand, and convenient interface. This allows casual users the ability to access the database without having to learn the query language, and it allows advanced users to quickly manage and manipulate the database without the trouble of entering formal commands using SQL. JDBA Tool also provides statistical data and information on who is using a database with its monitoring functions.

JServer Manager

JServer Manager is an intuitive graphical tool used to create, start, stop, backup, and restore databases. JServer Manager provides one central location for creating and managing all database servers at once.

JConfiguration Tool

JConfiguration Tool is a graphical tool for managing the configuration parameters for all databases. It provides a simple and direct method for modifying keywords in the DBMaker configuration files. Each of the configuration parameters is clearly defined within the user interface, eliminating the need to cross reference the documentation or memorize the definitions of keywords.

ESQL for C language

ESQL for C language is a graphical, interactive tool used to edit and preprocess Embedded SQL/C programs. It provides an easy-to-use interface for managing multiple ESQL/C programs, edit/preprocess them and examine the warnings/errors during preprocessing by simply clicking each of the error lines in the output window.

2.4 Syntax Diagrams

Syntax diagrams show the syntax for all SQL commands. These diagrams provide assistance when constructing a statement on the command line. An example syntax diagram is shown in Figure 2-1.

To use the syntax diagram, simply follow the line from start to finish. Any element of the command that cannot be bypassed is required. Any elements that can be bypassed are command options, and provide additional functionality for the command at the user's discretion.

Figure 2-1 Syntax of the ALTER TABLE Statement

Any words that appear in italics are placeholders for the actual names used in a database. The actual names should be substituted for these placeholders. In the above diagram, replace the <*table_name*> placeholder with the name of a table in the database. For example, in the tutorial database, replace the <*table_name*> placeholder with **Customers** to execute this command on the **Customers** table.

Also, note the direction of the arrows. Sometimes it is possible to have a list of items in a command, which is shown in the syntax diagram as a circular path. Both column name fields above can include a list of column names, separated by commas, as indicated by the circular path of the arrows.

3 System Architecture

This chapter explains the architecture of the two different models of DBMaker in detail. We will first look at the DBMaker process and the *Database Communication and Control Area (DCCA)*, which store all necessary information for each started database, and then go through the architecture of both models.

3.1 The DBMaker Process

A DBMaker process is a process that handles storage and retrieval of data according to user commands and other database functions. A DBMaker process consists of several layers as shown in Figure 3-1.

In this figure, we can see that the user applications communicate with DBMaker through an Application Program Interface (API). The API will pass user commands (SQL commands) or function calls to the *SQL Engine*, which is responsible for analyzing the SQL commands and translating them into sequences of function calls that are acceptable to the *Database Engine*. The SQL Engine passes these calls to the Database Engine, which executes these function calls to store data in tables or retrieve data from tables.

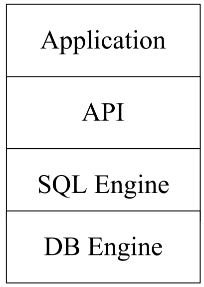

Figure 3-1: A DBMaker Process

The roles of the SQL Engine and the Database Engine are different. Basically, the SQL Engine handles SQL parsing and query optimization, while the Database Engine handles space/buffer management, concurrency control, crash recovery, and so on. All modules cooperate to maintain data consistency throughout the entire database. Most performance tuning parameters are related to the Database Engine.

The API and SQL Engines are the same in the single-user and client/server models. However, the Database Engines in the single-user and client/server models are different. The single-user model can handle only one user while the client/server model can handle multiple users.

In the client/server model the application and API are tied together, run on client machines while the SQL Engine and the Database Engine are tied together, and run on server machines. In this manner, the API can communicate with the SQL engine via network protocols.

3.2 Database Communication and Control Area (DCCA)

When started, DBMaker will first allocate a large block of memory to store database related information such as buffer pools and various types of control information. This memory block is called the Database Communication and Control Area (DCCA). It contains three types of data: *page buffers*, *journal buffers*, and the *System Control Area (SCA)*.

The DCCA is very important to the operation of DBMaker, especially when run in client/server mode. In Microsoft Windows and UNIX single-user environments, the DCCA is allocated from the private heap. In a UNIX client/server environment, the DCCA must be shared among all DBMaker processes that access the same database, so it cannot be allocated from the private heap. Instead, a shared memory mechanism, which is a standard function in UNIX, is used to allocate the DCCA. All DBMaker processes that run in client/server mode communicate with each other via the DCCA.

The size and the usage of the DCCA are easy to tune in DBMaker. This will affect the overall performance of DBMaker greatly. The DCCA is described in more detail in Chapter 17, "Performance Tuning".

3.3 Architecture of the Single-User Model

The DBMaker single-user model is a DBMS that supports only one user or application. It is smaller and faster than other models because it does not need to handle concurrency control. When one user or application owns a database, the single-user model of DBMaker is a good choice. Figure 3-2 shows the system architecture of the DBMaker single-user model.

Since only one user or application can connect to a single-user DBMaker database, the DCCA is obtained from the private heap and is not sharable. Note that there is no locking mechanism in the DBMaker single-user model. For the sake of performance, the DBMaker engine maintains all database data in memory while running, and writes the modified pages back to disk files, including data files and journal files, at the proper points in time. The **dmconfig.ini** file is a text file that defines many parameters required for DBMaker to configure itself.

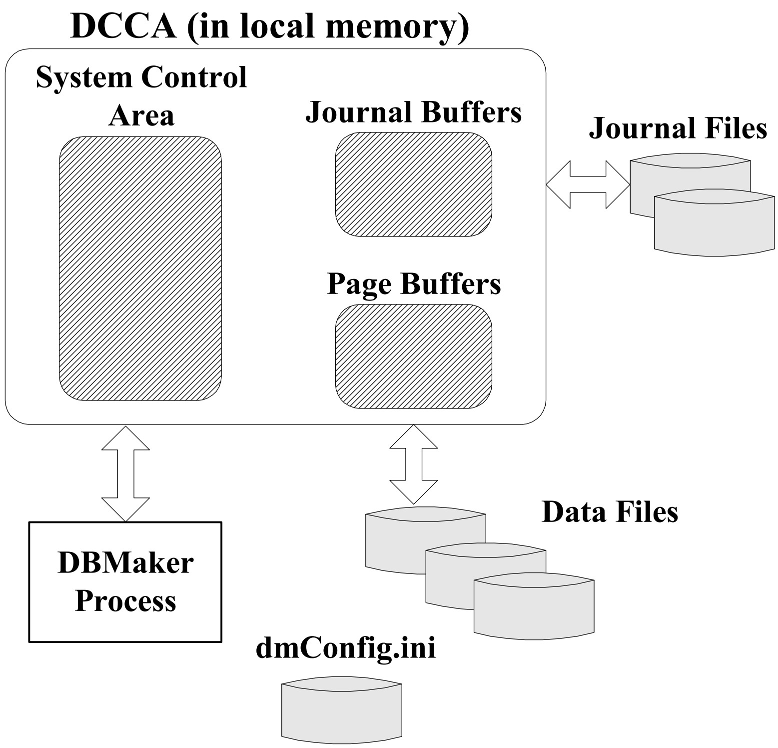

Figure 3-2: System architecture of the DBMaker single-user model

3.4 Architecture of the Client/Server Model

DBMaker may also run in the client/server model. In this model, there are two processes involved in an application program, the client application process, and the database server process. In general, the client process resides in a front-end PC or workstation and uses API library routines provided by DBMaker to communicate with the server process across a local area network. Note that in a client/server system, all involved machines including servers and clients can be of different platform types.

In the client/server version of DBMaker, a network management module is necessary both in the client and server ends. Network managers are responsible for sending data between the clients and the database servers. The network communications protocol is important in the client/server model. Currently DBMaker supports only one network protocol—TCP/IP (Transmission Control Protocol/Internet Protocol). If the client/server version of DBMaker is run on a system that does not normally support the TCP/IP protocol, it is necessary to install TCP/IP network software before running DBMaker. If the client application is run on UNIX, Windows 95, 98, XP, Windows NT, or Windows 2000, additional TCP/IP software does not need to be installed since all these operating systems have built-in support for TCP/IP. In Windows 95, NT, 98, 2000, or XP simply specify TCP/IP as one of the network protocols and install it on the system. Figure 3-3 shows the system architecture of the DBMaker client/server model.

On UNIX systems, when a client process connects to a database server, the DBMaker network server will fork another server process to handle subsequent queries. The original network server process will continue to wait for connections from other clients. Windows NT is a multithreaded system. The NT version of the DBMaker network server (**dmserver.exe**) is also a multithreaded program.

Therefore, when a client process connects to a server running on NT, the DBMaker server process will create another thread in its process space to handle the subsequent queries. The DCCA is allocated from local memory, not shared memory. There is always only one DBMaker server process for a database in Windows NT. There are more and more operating systems that support multithreading, DBMaker will use multithreading over process forking when possible. Current research indicates that multithreaded programs are more efficient than multi-process programs.

There are three components in the client/server model of DBMaker. These components are the server program, the client program, and the client library. The purpose of each of these components follows.

Server Program

The name of the DBMaker server program is **dmServer**. This program includes a network manager that deals with the network communication, and a database engine that handles the data access. This program must be started first so that client programs can connect to the database server.

Client Program

The name of the DBMaker SQL client program is **dmsqlc**. Use this program to connect to a database and issue SQL commands for data processing.

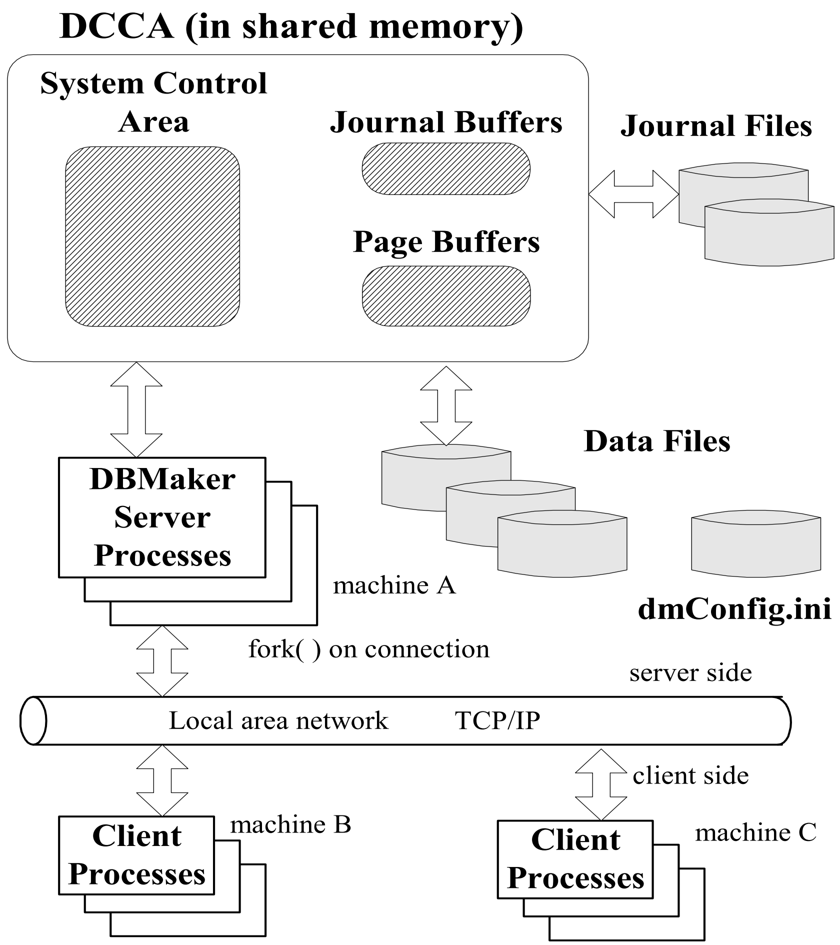

Figure 3-3: System architecture of the DBMaker client/server model

Client Library

The name of the DBMaker client library is **libdmapic.a** in UNIX, or **dmapi<*version number*>.lib** on Microsoft Windows systems. Users who intend to develop their own client programs must link their programs with the client library. For example, developers can use various development tools from many vendors to write their front-end applications. When building the applications, they must link those programs with the client library so that their custom applications can communicate with the server program.

4 Basic Database Administration

This chapter describes basic database administration, including creating a database, starting a database, connecting to a database, and shutting down a database. To perform the operations described in this chapter, database administrators can choose to use the command-line based dmSQL tool and edit the **dmconfig.ini** file, or use the JConfiguration Tool and JServer Manager utility.

The following sections describe configuration parameters and commands that are essential for basic database administration. The first section outlines the role and format of the configuration file. Subsequent sections describe the function of specific settings and how those settings affect database performance after it has been created, started, and connected to.

4.1 Configuration File - dmconfig.ini

DBMaker requires many configuration parameters whenever it is run. The DBMaker engine uses configuration parameters to specify how a database runs. File storage locations, runtime memory allocation, and network connections are just a few of the characteristics of a database that are set using configuration parameters. These parameters are stored as configuration variables in the **dmconfig.ini** file. A configuration variable is a keyword set equal to a value (refer to *Format* later in this section). Users can customize the database by setting parameters in the **dmconfig.ini** file or the JConfiguration Tool. The JConfiguration Tool simplifies management of configuration parameters through an easy to use graphical user interface. More information about JConfiguration Tool may be found in the "*JConfiguration Tool Reference*". Certain parameters (keywords) must be set before database creation, and others must be set before the database has been started. In addition, certain configuration parameters should not be changed after database creation or an error will be returned. The following sections describe how to manage settings by directly editing the keywords in the configuration file. To reference a complete list of **dmconfig.ini** options, see Appendix A at the end of this manual.

dmconfig.ini Location

DBMaker may search three places, for UNIX platforms, for the **dmconfig.ini** file:

- Current directory
- Directory specified by the DBMAKER environment variable
- The **data** subdirectory in the home directory for user **dbmaker** (**~dbmaker/data**)

DBMaker will search each location sequentially. If the relevant database section is not found in the **dmconfig.ini** file of one location, DBMaker will search in the next location.

However, for Microsoft Windows systems, including Windows 3.1, Windows 95, and Windows NT, the rule is different. DBMaker will only search for the **dmconfig.ini** file in the installed Windows system directory. In a typical Windows installation, this will be the **Windows** directory.

When DBMaker requires the value of a configuration parameter for a particular database, it will scan the above three directories (or the **Windows** directory on Microsoft Windows systems) to find a **dmconfig.ini** which contains a section having the same section name as the database. Use any text editor to edit this file and add or modify the parameter values in **dmconfig.ini** so that DBMaker will use them when it is running.

When a database administrator creates a database and no corresponding section can be found in any **dmconfig.ini** files, DBMaker will create a new section for the database in the first **dmconfig.ini** file found and use default values, or in a new **dmconfig.ini** file in the local directory (or the **Windows** directory on Microsoft Windows systems).

Therefore, when a database administrator starts a database the corresponding section in **dmconfig.ini** must be found, or DBMaker will return an error. Although various sections may be put in different **dmconfig.ini** files and different **dmconfig.ini** files put in different directories, this practice is not recommended. Using one global **dmconfig.ini** file will make maintenance easier.

JConfiguration Tool will display all database sections listed in the **dmconfig.ini** file. On a UNIX system, JConfiguration tool will display all sections of all **dmconfig.ini** files shown in the locations listed above.

dmconfig.ini Format

The **dmconfig.ini** file is divided into different sections. The first section lists the definitions of the most commonly used keywords. Subsequent sections begin with a header name that corresponds to the name of a database. The keywords under each section define the configuration of that database. Any string following a semi-colon is considered a comment.

➲ **Example**

A generalized format for a **dmconfig.ini** file:

```
[Section name1]
<key word1> = <value1>
<key word2> = <value2> <value3>; this is a comment
; this is a comment
...

[Section name2]
<key word3> = <value4> <value5>
<key word4> = <value6>
...
```

File Name and Size

A database consists of operating system files, and those files are defined in the **dmconfig.ini** file using keywords. The <*filename*> parameter is used in place of the <*value*> parameter. The **<*filename*>** parameter can be a simple file name like **firstdb.sdb**, a relative path like **mydb/firstdb.sdb,** or a full path like **/disk1/mydb/firstdb.sdb** ("/" for UNIX and "\" for Microsoft Windows).

The <*np*> parameter represents a number of pages. Pages are a unit of disk space allocated for a given file, and one page is roughly equivalent to four kilobytes.

➲ **Example**

A generalized format for indicating file names and sizes:

```
[Section name1]
<key_word1> = <filename>
<key_word2> = <filename> <filename>
<key_word1> = <np>
```

File Locations

If a database will be accessed by users who may run the DBMaker programs from different directories, (making the "current directory" different for each user), then all of the file names in **dmconfig.ini** should be full paths.

Alternatively, use the **DB_DbDir** configuration parameter. This keyword indicates the "home directory" (database directory) of a database:

➲ **Example 1**

The following sets the name of the database directory to db, instead of the default DB1 as indicated by the section header. Furthermore, other database files are placed in alternative locations, and on other disks.

```
 [DB1]
DB_DbDir = /disk1/db
DB_DbFil = mydb1
DB_JnFil = /disk2/usr/DB1.JNL
```

The resulting physical file names are:

```
DB DbFil -- /disk1/db/mydb1
DB JnFil -- /disk2/usr/DB1.JNL
DB_BbFil -- /disk1/db/DB1.SBB (using default file name)
```

➲ **Example 2**

Using the **DB_DbFil** keyword:

```
[DB2]
DB DbFil = mydb2
DB_JnFil = /disk2/usr/DB2.JNL
```

The resulting physical file names are:

```
DB DbFil -- mydb2 (in current directory)
DB JnFil -- /disk2/usr/DB2.JNL
DB_BbFil -- DB2.SBB (in current directory)
```

NOTE *The rule also applies to user-defined files.*

Some Important dmconfig.ini Keywords

The following list introduces some of the most important keywords with a short description of each. Keywords essential for database creation and startup are given in subsequent sections of this chapter. A complete list of keywords can be found in Appendix A. Examples of valid keywords that can appear in **dmconfig.ini**:

- **DB_DbDir**=<*filename*>—specifies the directory that the database files reside in
- **DB_DbFil**=<*filename*>—specifies the file name for the system database file as <***filename***>
- **DB_JnFil**=<*filename*>—specifies the file name for the system journal file as *<filename>*
- **DB_JnlSz**=<*np*>—specifies the size of the system journal file in <*np*> (number of pages)
- ***<logical_file>***=<*filename*> <*np*>—specifies that the user-defined file with the name <***logical_file***> will be mapped to <*filename*> with <*np*> pages. In other words, <*filename*> is the physical file name for <logical_file>

- **DB_NBufs**=<*np*>—specifies the runtime data buffer size in <*np*> (number of pages)
- **DB_SvAdr**=<*IP address*> or <*host name*>—specifies the database server's IP address or its host name. In a client/server system, this option must be set on the client side
- **DB_PtNum**=<*port number*>—specifies the TCP/IP port number used to communicate between the database client and database server
- **DB_MaxCo**=<*number*>— specifies the the maximum number of connections that the database can handle.

NOTE *Each DBMaker database page is 4KB. All pattern matching is case insensitive except for <logical_file>.*

Default Values

Some of the options have default values. Therefore, if a keyword does not appear in **dmconfig.ini**, its default value will be used. See Appendix A for a more detailed description of the keywords and their default values.

Sample dmconfig.ini file

In the following example, two sections are defined in the **dmconfig.ini** file, one for the **Personnel** database and the other for the **LIBRARY** database.

➲ **Example**

A typical **dmconfig.ini** file:

```
[Personnel]
DB_DbFil = /disk1/bin/PERSONNEL.DB
DB_JnFil = /disk1/bin/PERSONNEL.JNL
f1.os = /disk1/bin/PERSONNEL.OS 100
f1.blob = /disk1/bin/PERSONNEL.BLOB 1000
DB_UMode = 1           ; multi-user mode
DB_NBufs = 0           ; auto-configure number of data buffers
```

```
DB NJnlB = 100          ; number of journal buffers
DB MaxCo = 100          ; maximum number of connections
DB JnlSz = 2000         ; size of journal file (pages)
DB RTime = 0            ; restoration target time
DB SvAdr = 192.72.116.130       ; server's IP address
DB PtNum = 21000         ; and port number

[LIBRARY]
DB DbFil=/disk3/usr/lib/library.db
DB JnFil=/disk3/usr/lib/library.jnl
DB SvAdr = 192.72.116.137
DB PtNum = 26999
DB_JnlSz = 2000
```

4.2 Creating a Database

Creating a new database requires some planning. There are a number of configuration parameters that must be considered before creating a new database, some of which cannot be changed after the database has been created. Parameters that must be set during database creation are:

- Database name
- Database security (whether the database has different user authority levels)
- Case sensitivity (determines case sensitivity of certain schema objects within the database)
- BLOB frame size (the amount of disk space allocated for each BLOB frame)
- Language setting (determines the character set to be used- ASCII, Big5, etc.)
- Language code order (the pattern used to sort character type data)

All other configuration parameters can be changed after the database has been created, however, it is important to consider other database parameters before database creation. These parameters include:

- Tablespace name, location, size, and extensibility
- Number of journal files
- Journal file name, size, and location
- System data and BLOB files names, sizes, and locations
- Default user data and BLOB files names, sizes, and locations
- System temporary file name and location
- User-defined file names, sizes, and locations
- DBMaker log file locations
- Backup directory location
- Table replication log directory location
- Allow for user file objects
- Enable use of raw devices (for UNIX platform only)
- Enable client/server database
- Database IP address and port number (for client/server databases)
- Default user ID and password
- Memory allocation

DBMaker provides an easy to use wizard for creating databases in the JServer Manager tool. A database administrator can easily create a database by editing the **dmconfig.ini** file and using dmSQL. The following subsections outline the process for creating a database. The subsections also approximately follow the sequence of steps used in the JServer Manager create database wizard.

Naming the Database

Before naming a database, be aware of the following conventions:

- Database names can be up to 32 characters long

- Database names can contain any alphanumeric characters, including the underscore
- Character may be in any position
- Database names are not case-sensitive
- Database names must be unique among all computers that will connect to the database

Databases may be named from the create database wizard of the JServer Manager, or using dmSQL.

➲ **Example**

To create a database using dmSQL:

```
dmSQL> create db <database name>;
dmSQL> terminate db;
dmSQL> quit;
```

Schema Object Name Case Sensitivity

The case sensitivity of all identifiers in a database can be specified. Under the case insensitive mode, all identifiers appear in uppercase when defined. Once a database has been created, this setting cannot be changed. Setting the keyword equal to zero makes the database case sensitive. The keyword is set equal to one by default, so a database created without changing this setting will becase insensitive. The following **dmconfig.ini** variable specifies database case sensitivity:

DB_IDCap = <*value*> (default value = 1)

Setting Storage Parameters

There are ten different types of operating system files associated with a single database: system data and BLOB files, default user data and BLOB files, system journal files, a system temporary file, user-defined files, DBMaker log files, backup files, and a table replication log file. When a database is first created, the user may assign names and locations for each file, or DBMaker can assign default values to them. It is important before creating a database to have a good understanding of the general role these files serve within a database.

Many of the parameters discussed in this section may be modified from the Storage page of the JConfiguration Tool. To learn more about how to use JConfiguration Tool to change database parameters see the “*JConfiguration Tool Reference*”. More information about managing files is available in section 3.2 *File Types*.

When creating a database, DBMaker will create the system database file, the journal file, and the system BLOB file according to the related settings in the **dmconfig.ini** file. If no **DB_DbFil**, **DB_JnFil**, or **DB_BbFil** settings are defined, the default setting will be used.

The default values are:

DB_DbFil -- database name + '.SDB'

DB_JnFil -- database name + '.JNL'

DB_BbFil -- database name + '.SBB'

SPECIFYING THE DATABASE DIRECTORY

The database directory is the default location where files associated with a database are created and stored. If the defined file is specified with a full path name, DBMaker will use that name to reference it. If a file name without a full path is used, DBMaker will search for the database directory. If it is not found, DBMaker will use the file name and assume it is located in the current directory. The following keyword is used in **dmconfig.ini** to specify the database directory:

When creating a new database in Windows, DBMaker assigns a default database directory of (DBMaker Installation Directory)/bin. You must create a new directory for the database files to reside in. Multiple databases must not be created in the same database directory.

DB_DbDir = *<pathname>* (default: *<DBMaker installation directory>/bin/<Database Name>)*

➲ **Example**

To set the database directory to **/disk1/db**:

```
[DB1]
DB_DbDir = /disk1/db
```

CREATING THE SYSTEM TABLESPACE

A DBMaker database is composed of several logical divisions known as tablespaces. With tablespaces, the database can be divided into manageable areas. In the logical view, a tablespace contains one or more tables and indexes. In the physical view, a tablespace is the physical storage that consists of one or more files. A newly created database will have two tablespaces, the system tablespace, and the default user tablespace.

The *system tablespace* consists of a system data file and a system BLOB file. The system tablespace is used to record the system catalog for the entire database. The database administrator may specify the initial location of system data and BLOB files in the system tablespace.

The system tablespace cannot be dropped (deleted), although other user tablespaces can be. The initial size of the system database file is 150 pages (600KB). The following keywords in **dmconfig.ini** are used to define the system tablespace:

System data file: **DB_DbFil** = <*filename*> (default: " <*Database Name*>*.SDB*")

System BLOB file: **DB_BbFil** = <*filename*> (default: " <*Database Name*>*.SBB*")

The ***<filename>*** parameter can be a simple file name like **firstdb.sdb**, a relative path like **mydb/firstdb.sdb**, or a full path like **/disk1/mydb/firstdb.sdb** ("/" for UNIX and "\" for Microsoft Windows).

➲ **Example**

Entering the following lines into the dmconfig.ini file will result in the system tablespace files being stored in the **/disk1/mydb/** directory.

```
DB_DbFil = /disk1/mydb/firstdb.sdb
DB_DbFil = /disk1/mydb/firstdb.sbb
```

Creating the User Default Tablespace

The *default user tablespace* initially contains one data file and one BLOB file. User data is stored in these files. The database administrator may specify the initial size and location of data and BLOB files in the user default tablespace. Data file size is specified in pages (1 page = 4K). BLOB file size is specified in frames. Frame size can be defined by the user and is discussed in "Related Topics:

Specifying the BLOB Frame Size" later in this chapter. By default, the user default tablespace is autoextend. This means that if the tablespace is full of data, DBMaker will enlarge the files (and therefore the tablespace) automatically. However, it is more flexible and efficient to create additional tablespaces to store user tables.

JDBA Tool can help to create new tablespaces and manage existing ones. If data or BLOB files are added to a tablespace without specifying a full path, the file will be created in the database directory. The user default tablespace cannot be dropped (deleted), although other user tablespaces can be. The following keywords in **dmconfig.ini** are used to define the default user data and BLOB files:

User data file: **DB_UsrDb** = <*filename*> (default: " <*Database Name*>*.DB*")

User BLOB file: **DB_UsrBb** = <*filename*> (default: " <*Database Name*>*.BB*")

The **<*filename*>** parameter can be a simple file name like **firstdb.sdb**, a relative path like **mydb/firstdb.sdb**, or a full path like **/disk1/mydb/firstdb.sdb** ("/" for UNIX and "\" for Microsoft Windows).

Creating Journal Files

Journal files provide a real-time, historical record of all changes made to a database, and the status of each change. Up to eight journal files can be created. Each journal file has a fixed size. When all journal files are filled by active transactions (i.e. transactions are not committed, and their occupied journal blocks cannot be freed), the current transaction will be aborted because no space is available; this is called journal full. Make sure that the longest transaction will not use all journal records in all the journal files. If journal files are created without specifying a full path, then they will be created in the database directory. Journal files may not be modified after the database has been started. To reduce the number of journal files, add more journal files, or change journal file size, restart the database in new journal mode. More information about new journal mode is available in section 4.4, “Starting a Database”. Also, refer to section 5.2, *File Types* for more information on journal files. The following keywords in **dmconfig.ini** are used to define journal file names, locations, and sizes:

Journal file name(s): **DB_JnFil** = <*filename*> (default: "<*Database Name*>.*JNL*")

Journal file size (pages) **DB_JnlSz** = <*np*> where *np*=100~524287 pages (default: 1000 pages)

➲ **Example**

The following lines in the **dmconfig.ini** file tell DBMaker to create two 500 page journal files on two separate disks in the **/mydb** directory

```
DB_JnFil = /disk1/mydb/firstdb1.jnl /disk2/mydb/firstdb2.jnl
DB_JnlSz = 500
```

CREATING SYSTEM TEMPORARY FILES

System temporary files are used by DBMaker to store information about the database, such as sorting results, while the database is active. These files are generated when it is necessary and deleted when the database is shut down. If temporary files are created without specifying a full path, then they will be created in the database directory. Up to eight system temporary files may be specified. Each temporary file may hold up to 2 gigabytes. Temporary files may be located on different disks for improved disk I/O performance. Users should reserve enough space on disk for an entire temporary file (maximum of 2GB for a single file), or errors may be returned. System temporary files may be specified either by using JConfiguration Tool or by editing **dmconfig.ini** before starting the database. The following keyword in **dmconfig.ini** is used to define system temporary file names and locations:

DB_TpFil = <*filename*>[<*filename*>…] (default: "<*Database Name*>.*TMP*")

SPECIFYING THE BLOB FRAME SIZE

A BLOB frame is the smallest unit of storage used by BLOB files, which are used to store large object data such as LONG VARCHAR or LONG VARBINARY. BLOB frame size cannot be changed after creating the database. The minimum frame size is 8KB and the maximum frame size is 256KB. Determining the frame size is a trade-off between disk utilization and performance. If entire BLOBs are frequently retrieved, adjusting the frame size to contain an entire BLOB will result in better performance because only one disk access is required. However, there may be large variations in the size of the BLOB data. If the frame size is large enough to contain the largest BLOB, it could waste disk space, as other frames that contain smaller BLOBs will contain unused disk space. Alternatively, if frames are only large enough to contain the smallest BLOBs, performance will be degraded when fetching larger BLOBS that are stored in multiple frames. The following keyword in **dmconfig.ini** is used to specify BLOB frame size:

DB_BFrSz = <*nk*>. The < *nk* > parameter is the frame size in kilobytes. The size of the system BLOB file is (4 + (number of frames – 1) × *nk*). Refer to Chapter 7, *Large Object Management* for more detailed information.

➲ **Example**

To set the BLOB frame size to 10 KB:

```
DB_BFrSz = 10
```

Setting the Number of Pages to Extend an Autoextend Tablespace

When all pages in the data file or BLOB file of an autoextend tablespace are full, DBMaker will automatically extend the number of pages or frames in the file to allow the tablespace to grow. This setting tells DBMaker how many pages or frames to add to the full file in the event that it is filled. If the Database Administrator expects that the database will grow very quickly, then a higher number should be selected to lessen the frequency at which the file is appended. This number can be adjusted before starting a database by using JConfiguration Tool, or by editing the **dmconfig.ini** file. The following keyword in **dmconfig.ini** is used to specify the number of pages/frames to extend an autoextend tablespace:

DB_ExtNp = <*np*>, where <*np*> is the number of pages to extend (default: 20 pages / frames)

ENABLING USER FILE OBJECTS

FILE type data can be stored as user file objects or system file objects. User file objects are external files that are accessed through the machine where the database resides. In other words, a user file object is only a link to an external file outside the database. Enabling user file objects allows a FILE type column to link to the external files, which will be accessed by the database server. It may be disabled and enabled as needed. Inserted user file objects can be accessed even if the setting is turned off. User file objects may be enabled through the storage page of the JConfiguration Tool before starting the database. The keyword value may be modified before starting the database. Setting the keyword equal to 0 prevents file objects from being inserted. Setting the keyword equal to 1 allows file objects to be inserted. The following **dmconfig.ini** variable toggles file object capability:

DB_UsrFo = <*value*> (default: 0 / disabled)

CREATING A DIRECTORY FOR SYSTEM FILE OBJECTS

System file objects are created, deleted, and managed by DBMaker. All system file objects are placed in subdirectories of the system file object directory. Changing the system file object directory does not change the location where previously inserted system file objects reside. The system file object name and location may be set from the storage page of the JConfiguration Tool before starting the database, or during runtime with JServer Manager or JDBA Tool Run Time settings. The keyword value may be modified before starting the database. The following **dmconfig.ini** variable sets the name and location of system file objects:

DB_FoDir = <*pathname*> (default: "\<*database directory*>*fo*")

CREATING A DIRECTORY FOR USER-DEFINED FUNCTION DLL FILES

The database administrator may specify the directory where the dynamic link libraries (DLL) of user-defined functions (UDF) are placed. UDFs are compiled functions stored in a dynamic link library (.DLL for Windows operating system, or .so for UNIX operating system) that the user wants to be able to use in DBMaker. The DLLs stored in the Directory of User-defined Function DLL files are accessible to DBMaker and can be used in SQL statements or ODBC applications. UDFs should be loaded when the database starts. The following keyword specifies the location of UDF DLL files:

DB_LbDir = <*filename*> (default: current working directory)

Raw Devices

The DBMaker physical storage system is very flexible. In a UNIX system, DBMaker allows users to create a database with UNIX files only, with raw device files only, or with files from both file systems. In **dmconfig.ini**, if a file name begins with ***/dev/***, that file will be treated as a raw device.

I/O operations on raw devices will be faster than on regular UNIX files, so database administrators are encouraged to use raw devices as database files. To use raw devices as database files, the system manager must create raw devices before creating any databases. Please refer to the UNIX system manual for the procedure to create raw devices.

Multiple files can be put on a raw device without partitioning the raw device. To put multiple files on a raw device you must consider the following constraints:

- When you set multiple files on a single raw device, the files cannot be autoextend files.
- When setting multiple files on a raw device the file size cannot be changed after the initial set up.
- The total size of all files in a single raw device is restricted to 8TB.
- If an autoextend file is placed on a raw device no other files can be put on the device. Other than the files you set as autoextend, the DB_DBFIL, DB_BBFIL, DB_USRDB, DB_USRBB, and DB_TPFIL files are all autoextend files.
- If DB_DBFIL, DB_BBFIL, DB_USRDB, DB_USRBB, and DB_TPFIL are set to a raw device, they can have only one parameter; number of pages. You cannot set an offset to them. For example :

This is valid. It will create a file with 500 pages.

```
DB_DBFIL = /dev/sda 500 ;
```

This is valid too, but it will create a file with 30 pages. The parameter 500 will be ignored.

```
DB_BBFIL = /dev/sdb 30 500 ;
```

NOTE *Microsoft Windows and Windows NT do not support raw devices.*

➲ **Example 1:**

```
[MYDB]
f1 = /dev/sda 0 500
f2 = /dev/sda 500 200
f3 = /dev/sdb 300
```

To create a regular tablespace, ts1, containing the above raw device files:

```
DmSQL>CREATE TABLESPACE ts1 DATAFILE f1, f2, f3
TYPE=DATA
```

Then you will create three files in the raw devices. The first file has the size 500*4k = 2000k starting at address 0 in /dev/sda, and the second has the size 200*4k = 800k starting at address 500*4k = 2000k in /dev/sda. The third has the size 300*4k =1200k starting at address 0.

➲ **Example 2**

```
[MYDB2]
DB JnlSz = 1000
DB_JnFil = J.1 /dev/sda 1000 /dev/sda 2000 J.2jnl
```

Then you will create two normal journal files J1.jnl, J2.jnl and two raw device journal files that one starts at the address 4000k of the /dev/sda and the other at the address 8000k of the /dev/sda.

Enabling Client/Server Database

Any database can be started as a single-user database or a multi-user database. Before creating the database, determine what the primary function of the database is and which user mode is more suitable. If the database is to be primarily a multi-user database, configure the database to use an IP address or DNS name that is appropriate for the network that the DBMaker server will be running on. Also, specify the TCP/IP port number that the database server will use. The client side database will also use this information to connect to the database. This setting can be changed any time before starting the database, but we highly recommend setting these parameters before database creation to ensure smooth operation. Clients will be unable to connect to an improperly configured server database. If both settings are disabled, the database will start in single-user mode. These parameters can be altered from the connections page of the JConfiguration Tool, or by editing the following **dmconfig.ini** keywords:

IP address/Server name: **DB_SvAdr** = <***IP_address***> or <***host name***> (default: local host name or 127.0.0.1)

Port number: **DB_PtNum** = <*port number*> (default: 2300, 1024~65535)

Default User and Password

The default user name and password must already exist in the database. These two keywords are not examined when starting a database, but are checked when connecting to a database instead.

➲ **Example**

To specify a default user name and password to use when connecting to a database:

```
DB_UsrId = <user name>
DB_PasWd = <*****>
```

Changing the Language Code Order

The sort order definition file will affect the result of indexing, sort ordering, predicate, etc., for CHAR and VARCHAR data. Other data types are not affected.

DBMaker supports different character sets (language codes), such as US-ASCII for English, BIG5 for traditional Chinese, GBK for simplified Chinese, and JIS for Japanese. The keyword DB_LCode in **dmconfig.ini** file specifies which character set DBMaker will use. For each character set, there may be several sort orders.

In traditional Chinese for example, the sort order may be according to code sequence, stroke count, or phonetic equivalent. The default sort order for DBMaker is binary sequence. While creating a new database, the user-defined order definition file specified by the keyword DB_Order could change the behavior of the sort order. The language code parameter may also be set from the Create Database page of the JConfiguration Tool.

Setting Sort Order of Database

The following example shows how to set the local language and the sort order file before a new database is created.

➲ **Example**

To set the language type to traditional Chinese, use BIG5:

```
[MY DB]
………
DB LCode =1                ; BIG5 for traditional Chinese
DB_Order = big5_stroke.ord ; order definition file
```

The keyword DB_Order indicates the user-defined order definition file named **big5_stroke.ord**, which should be placed in the **shared/codeorder** subdirectory of the DBMaker installation directory. The order definition file is a pure text file, which affects the sorting result in DBMaker. The keyword is used when the database is created and then it is recorded in the database and not used. Without this keyword, while creating the database, the sorting sequence would be in a binary sequence. Once a definition file has been specified, it must always exist or the database will fail to start.

User-defined Order Definition File

The order definition is a user-defined pure text file. The order definition file arranges the sequence of valid characters. An example of the naming scheme looks like *codename_ordertype.ord*, where *codename* is the name of language code and *ordertype* is the type of ordering e.g. *big5_stroke.ord***.**

Example

An order definition file:

```
   Comment: Write information here.

[BEGIN]       // begin to arrange the character sequence

c             // ASCII 0x63
0x62          // Character 'b'
a             // ASCII 0x61

[SINGLE]      // Single-Byte Character Default Order

[DOUBLE]      // Double-Byte Character Default Order

0xA440        // one of Chinese characters
0xA441        // one of Chinese characters
0xA442        // one of Chinese characters
```

All lines before the [**BEGIN**] keyword are regarded as comments. All words after // or /* are also comments. After the [**BEGIN**], each line represents one character and should occupy the first position of the line followed by at least one space or a new line of characters. In the above example the character **c** is less than **b** and **b** is less than **a**.

If the text editor cannot be used to edit some characters, represent them with hexadecimal. For example, character **a** can be written as **a** or its code value **0x61**. It is also very useful for invisible characters.

The creator of the sort order may only be interested in some characters and let others be sorted by default, i.e. binary. The keywords [**SINGLE**] and [**DOUBLE**] can be used to represent the single character set and double character set, both of which are not specified in the definition file. If there is not a [**SINGLE**], the absent single-byte characters will come before all characters in the definition file. If the [**DOUBLE**] is absent, the absent double-byte characters will come after characters in the definition file.

DBMaker will ignore all errors found in the definition file. For example, if [**BEGIN**] is lost, DBMaker will always use the default sorting order for all characters. If the same character appears two or more times, only the first one will be processed; other characters will be ignored. The database administrator should check the sort order behavior carefully to see if it is correct after creating a database.

In distributed database environments, all databases should use the same sort order definition file. When copying or moving a database to another machine, do not forget to copy any existing sort order definition files.

The Data Communications and Control Area

The *Data Communications and Control Area (DCCA)* is a memory block in which almost all information and data is placed. For multi-user databases, the DCCA is allocated from shared memory and is used to do inter-process communications. When a database starts, it will allocate a DCCA to hold all information about that database. The DCCA can be divided into three parts—page buffers, journal buffers, and the system control area.

There are several keywords in **dmconfig.ini** related to the usage of the DCCA:

- **DB_NBufs**=<*np*>—this keyword sets the number of page buffers which DBMaker will use. The default value is 0 (automatically configure).
- **DB_NJnlB**=<*np*>—this keyword sets the number of journal buffers which DBMaker will use. The default value is 64.
- **DB_ScaSz**=<*np*>—this keyword will set the number of pages in the system control area. The default value is 200.
- **DB_MaxCo**=<*number*>— specifies the maximum number of concurrent transactions that the database can handle. **DB_MaxCo** is also used for formatting the journal file when the database is created or started with a new journal.

The size of the DCCA can be estimated by adding the size of the page buffers, the journal buffers, and the system control area. When the specified size of the DCCA is not large enough, DBMaker will automatically allocate the minimum necessary space to hold the information required to the DCCA instead of the default size used in the calculation above.

The size of the DCCA cannot exceed the allowable shared memory size of the system in a multi-user environment in UNIX, because in such a case the DCCA is allocated from shared memory. Users can refer to their UNIX manuals for instructions on how to increase the size of shared memory, which generally requires a rebuild of the kernel. DBMaker will run more smoothly with more buffers and a larger system control area.

The relationship between the DCCA, page buffers, journal buffers, and the system control area is explained in more detail in Chapter 17, "Performance Tuning".

DCCA parameters may also be set from the Cache and Control page of the JConfiguration Tool. For details, refer to the *JConfiguration Tool Reference*.

4.3 Starting a Database

The purpose of starting a database is to allocate the required resources from the operating system, initialize them, and wait for users to connect. The settings of certain configuration parameters must be considered before starting a database. These parameters include:

- Database startup mode
- Enable client/server database
- Database IP address and port number (for client/server databases)
- Default user ID and password
- Memory allocation
- Method for reporting errors

A database may be started using dmSQL or JServer Manager. For more information on starting a database using dmSQL, refer to the following sections. To find out how to use JServer Manager to start a database, refer to the "*JServer Manager User's Guide*".

Single-User

A user must start a single-user database every time they want to connect, and terminate the database when they finish using it.

➲ **Example :**

To start a single-user database using dmSQL:

```
dmSQL> START DB <database name> <user name> <password>;
.
< do DML here >
.
dmSQL> TERMINATE DB;
```

NOTE *Only users with DBA privilege can start a database. For information about database privileges, refer to Chapter 9, "Security Management". If a database is started in single-user mode, only one user can access the database at a time.*

Client/Server

The DBA must start the client/server database on the server machine so that all clients on remote machines (or on the same machine) can connect to the server database via the network. Two configuration variables must be set on the server before the database is started.

Starting a client/server database is a little more complicated than starting a single-user database. First, we need to know the server machine's IP address. The only ID to distinguish each machine on a network is the IP address. The **dmconfig.ini** keyword **DB_SvAdr** specifies the server's IP address.

The second item is the port number. The server program will bind to a given port number, specified by **DB_PtNum** in **dmconfig.ini**, to wait for connections. All client programs must connect to that port number in order to communicate with the database server.

Example 1

To specify the Server IP address and the Server and Client port numbers:

```
DB SvAdr = <server IP address> (on client side)
DB_PtNum = <port number> (on both server and client sides)
```

Example 2

To start a client/server database on the server machine using dmServer:

```
UNIX> dmserver <database name>
```

Example 3

Enter the user name and password. dmServer will start the database and wait for clients to connect:

```
UNIX> dmserver [-f] [-t port number] [-u username [-p password]]
      database_name
```

Description of Unix switches:

- **f** — Runs the server program in foreground mode. dmServer normally runs in background mode.
- **t** — Specifies the port number to use. This port number will be used rather than the port number defined in **dmconfig.ini**.
- **u** — Specifies the login user name.
- **p** — Specifies the password for the given user name.

If a username and password are not specified on the command line, dmServer will search for the **DB_UsrId** and **DB_PasWd** in **dmconfig.ini**. If not found, dmServer will prompt users to enter a username and password.

Start Mode

Specify the start mode of a database by using the **DB_SMode** keyword in **dmconfig.ini**. The **DB_SMode** keyword may have six values, corresponding to six start-up modes:

- **1** — Normal start starts up a system normally. If the database crashed in the last session, DBMaker will perform crash recovery automatically to bring the database to a consistent and stable state.
- **2** — New Journal. The database should be set to start in new journal mode if new journal file names and/or locations have been set in the **dmconfig.ini** file. New journal file names and locations may also be specified on the Storage page of the JConfiguration Tool. All old records will be overwritten if the previous journal file names are kept. This setting must be selected if the user wants to change the journal file size, add more journal files, or change the journal file name. It is recommended to do an incremental or full backup before selecting this option.
- **3** — Restore Backup Database uses the backed up database files (including the journal file) to start the database. DBMaker will use the incremental backup files to roll over the operations up to the point in time specified by **DB_RTime**. If no value is specified or the date specified is later than the time of the last incremental backup, **DB_RTime** will revert to its default value. Refer to Chapter 14 *Database Recovery, Backup, and Restoration* for more detailed information on rollover.
- **4** — Source for Target Database is used for database replication. Starting up a system in this mode makes it a primary (source) database. For more information on database replication, refer to Chapter 16, *Data Replication*.
- **5** — Target of Database Replication is used for database replication. Starting up a system in this mode makes it a slave database. For more information on database replication, refer to Chapter 16, *Data Replication*.
- **6** — Database is Read Only starts up a system normally, but the database is read-only or only provides read privilege to all users. Starting a primary database in read-only mode prohibits users from modifying it.

Start mode may also be specified on the Start Database page in JConfiguration Tool or the Start Database Advanced Settings window in JServer Manager.

Forced Startup

When attempting to start a damaged database, it is possible that an error message will always be returned. The only solution is to use "Forced Startup" provided by DBMaker. Set the configuration variable **DB_ForcS** to one and DBMaker will force the database to start up. Refer to Chapter 14, *Database Recovery, Backup and Restoration* for more detailed information.

E-mail Error Report System

Typically all error messages are stored in the error.log file. Unless the database administrator consistently checks the error.log file, certain database errors may pass unnoticed. DBMaker provides an e-mail error report system to ensure that database administrators are made aware of errors in the system.

The error report system may be activated either by setting two configuration file keywords, with the JConfiguration Tool, or during database startup with the JServer Manager. The keywords that govern the behavior of the e-mail report system are **DB_ERMRv** and **DB_ERMSv**. Use **DB_ERMRv** to specify the recipient addresses for error report e-mail, and use **DB_ERMSv** to set the address of an SMTP server to route e-mail through. For more information on setting the e-mail error report system with JConfiguration Tool or JServer Manager, refer to the *JConfiguration Tool Reference* and the *JServer Manger User's Guide*, respectively.

4.4 Connecting to a Database

This section discusses how to connect to a running client/server database. A user must first connect to a database before performing DML operations. After disconnecting, a client/server database is still active. Users can continue to make connections until the database is shut down.

Certain parameters exist for client/server connections, including port number, server address, connection time-out interval, and lock time-out interval. Connection parameters are set by changing keyword values in the **dmconfig.ini** file or by using the JConfiguration Tool.

A single-user database only allows a single user connection, every time a user is going to use the database they must start it, they do not need to make a connection. See “Starting a Database” for more information.

Client/Server Database

The **DB_SvAdr** and **DB_PtNum** keywords must be set in the **dmconfig.ini** file. If the **DB_UsrId** and **DB_PasWd** keywords are defined in **dmconfig.ini**, the **<*username*>** and **<*password*>** options in the CONNECT command can be ignored.

➲ **Example**

To connect to and disconnect from a client/server database with dmsqlc:

```
dmSQL> CONNECT TO <database name> <username> <password>;
.
< do DML here >
.
dmSQL> DISCONNECT;
dmSQL> QUIT;
```

Connection Time-Out

In a client/server model database, sometimes a client cannot connect to the server because the server machine is powered off or the IP address of the server machine is wrong. In these cases, users must wait for the connection to be established. To set the connection time-out value, users can set the **DB_CTimO** parameter (in seconds). The default value for this keyword is 5 seconds.

Lock Time-Out

Locks are required for concurrency control between multiple transactions on the same database objects. For more information on transactions and concurrency control, refer to Chapter 9, *Concurrency Control*. When connecting to a database a lock time-out keyword, **DB_LTimO**, should be defined in the **dmconfig.ini** file to indicate how long (in seconds) a user will wait for a lock that cannot be acquired.

For example, if **DB_LTimO**=10, DBMaker will return a "lock time-out" error if the user waits for a lock for more than 10 seconds. **DB_LTimO** can be set to zero indicating that user does not want to wait at all. Setting **DB_LTimO** to -1 will turn off this feature. In this case, a user will wait for a lock until the lock is released. Each user can have a **DB_LTimO** value.

4.5 Shutting down a Database

A database should be shut down after all operations are finished. DBMaker will free all resources, such as the DCCA, for the operating system. If there are still active transactions in the database engine, DBMaker will abort them.

However, if there are still active connections to the database engine, DBMaker will shut down the database without killing the processes for those connections. In this case, the database administrator should manually kill the processes; otherwise, the error message "Cannot lock file transaction rollback" will occur when starting the database the next time.

Therefore, database administrators (DBA users) should ensure that all users are logged off before shutting down the database. To shut down a database, a DBA has to connect first and then issue the proper command. Only a DBA has the privilege to shut down a database.

➲ **Example**

To shut down single-user or client/server databases, use dmSQL:

```
dmSQL> CONNECT TO <database name> <DBA username> <password>;
dmSQL> TERMINATE DB;
dmSQL> QUIT;
```

5 Storage Architecture

This chapter introduces the storage architecture of DBMaker. The storage architecture of DBMaker includes the *logical level* and the *physical level*.

The logical level is the view that is presented to users, and organizes data in the database in a way which is easy to understand. The physical level consists of operating system files which correspond to information in the tablespaces, but which are managed by DBMaker and hidden from the user.

This chapter also explains how to control the storage allocation of a database by using tablespaces and files.

5.1 Architecture

A DBMaker database is composed of one or more logical divisions known as *tablespaces*. Tablespaces are the primary logical storage structure in DBMaker. In the logical view, a tablespace contains one or more tables and indexes (see Figure 5-1). In the physical view, a tablespace is the logical storage that consists of one or more operating system files (see Figure 5-2).

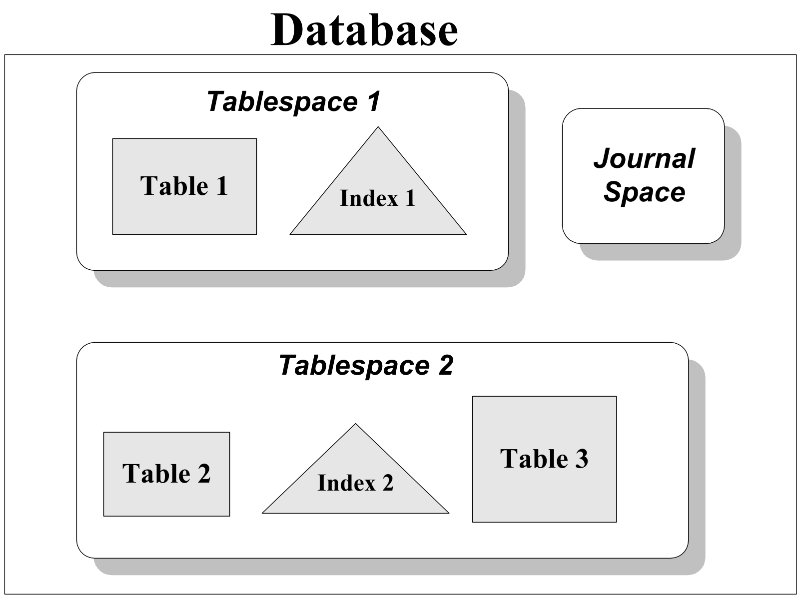

Figure 5-1: DBMaker database storage components in the logical view

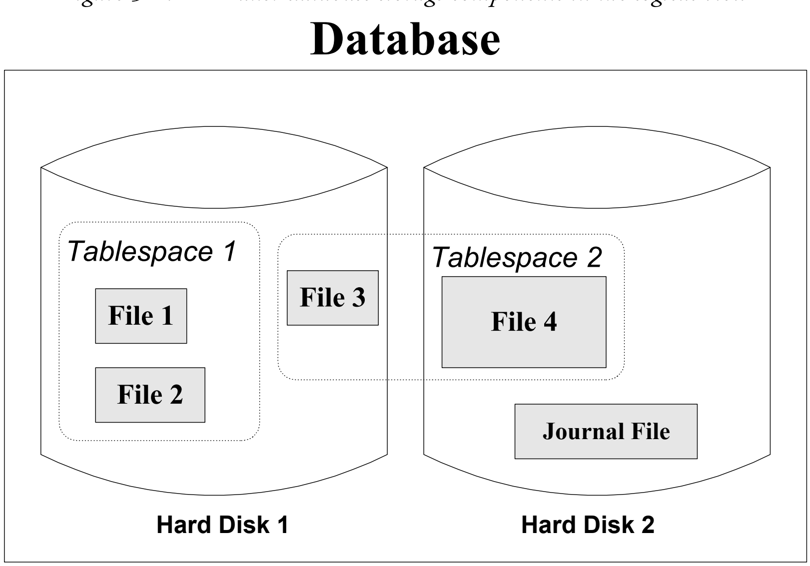

Figure 5-2: DBMaker database storage components in the physical view

5.2 File Types

Ten different operating system file types are used in DBMaker to store different aspects of a database: system data and system BLOB files, user data and user BLOB files, system journal files, a system temporary file, user-defined files, DBMaker log files, backup files, and a table replication log file. The system data file, the system BLOB file, user data files, and user BLOB files are of primary concern regarding database storage architecture and tablespaces. Journal files play an important role in storing records of transactions performed on the database, and are vital to database backup and recovery.

To increase database performance, DBMaker places data into two different types of files—data files and *Binary Large Object (BLOB)* files. BLOB data consists of large data objects in the form of image, voice, or large text, which cannot be packed into a page. DBMaker stores the BLOB data in BLOB files and stores the data rows and index keys in the data files. In order to achieve high performance, DBMaker manages these two file types in different ways.

User Data Files

Data files are comprised of pages, while BLOB files are comprised of frames. The maximum size of both data and BLOB files is 8TB. However, there are two major differences between frames and pages:

- The size of a page is fixed at 4KB, but a user can customize the size of a frame.
- A page can contain more than one tuple, but a frame only contains a single BLOB data item.

A data page, 4096 bytes, is the smallest unit of storage used by data files. The data page format is similar regardless of whether the data page stores table or index data. A data page contains four sections: the page header, row data, free space, and the row directory.

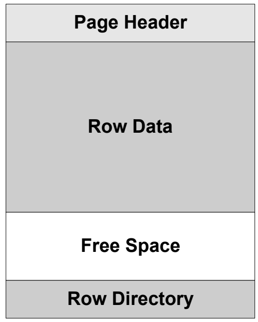

Figure 5-3: Format of a data page

The page header contains general page information for the DBMaker system. The row data area contains the actual table or index data that is displayed as rows and columns when looking in a table or index, and the row directory contains information about the rows in the page. Free space is the available space on that page that has not yet been used to store data.

User BLOB Files

A BLOB frame is the smallest unit of storage used by BLOB files. The size of the BLOB frame can only be set to a value other than the default before creating a database. The minimum frame size is 8KB and the maximum frame size is 256KB. A BLOB frame contains three sections: the frame header, BLOB data, and free space. For more information about BLOB files, refer to Chapter 7, *Large Object Management*.

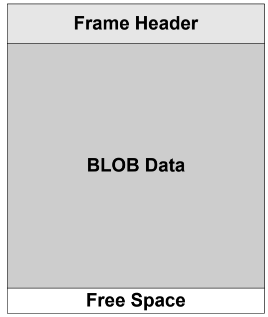

Figure 6-4: Format of a frame.

Like the page header, the frame header contains general frame information for the DBMaker system. The BLOB data area contains the BLOB data, and each frame can only contain a single BLOB item. However, BLOB data that is larger than the frame size can be spread over several frames. Free space is the available space on that page that has not been used to store BLOB data.

Journal Files

DBMaker's journal is composed of one or several physical journal files. Every journal file is a fixed-size file, and all journal files are the same size. Journal file size is determined in pages (4096 bytes per page) via the configuration file. Internally, a journal file is composed of blocks, where each block is 512 bytes. Every action that causes a change in the database system will have a journal record to record it. Journal records are the logical elements in the journal, and several journal records may be packed into a journal block or a single journal record may span several blocks. A journal record that is owned by an active transaction cannot be reused.

All journal files form a ring of journal records; journal records are written to sequential journal blocks from the beginning of the file to the end. If the database has been configured to have more than one journal file, DBMaker automatically switches to a new journal file when the current file fills. Otherwise, journal records will be written over journal blocks at the beginning of the journal file. When all journal files are filled by active transactions, the current transaction will be aborted because no blocks are available; this is called *journal full*.

In addition to journal records, a journal file contains some blocks to record the journal status, called journal status blocks. These are used when recovering or restoring the database. Recovery and restoration will be described in later sections.

DBMaker maintains journal block buffers in memory to speed up file access. Before the actual modified data is written to disk, the journal record is written to disk using the *Write-Ahead-Log (WAL)* protocol. When the journal buffer is full or a transaction is committed, the buffer will be flushed to the journal files in accordance with the WAL protocol.

Journal Parameters in DMCONFIG.INI

Several journal file parameters can be set to enhance database performance.

- **DB_JnFil**—Specifies the names of journal files. One to eight journal file names can be specified. A comma or a space separates every journal file name.

➲ **Example**

The database will have seven journal files specified on different drives to enhance performance:

```
DB JnFil=myDb.jn1, myDb.jn2, myDb.jn3, /disk1/usr/myDb.jn4,
         myDb.jn5, /disk2/usr/myDb.jn6, myDb.jn7
```

- **DB_JnlSz**—Specifies the size of a journal file as a multiple of journal pages. (One journal page is 4096 bytes.) The total journal file size is:

  (number of journal files × journal file size) pages

Decide on a reasonable size for journal files when creating a database. As the previous section stated, when all journal files are filled, the current transaction might be aborted because of a full journal. Therefore, a small journal file size may cause a long transaction to be aborted by the system. If database operations involve long transactions, choose a larger journal file size or more journal files.

- **DB_NJnlB**—Specifies the size of a journal buffer as a multiple of journal pages. (One journal page is 4096 bytes.)

Resizing Journal Space

If journal full messages are frequently encountered when a database is running, enlarging the journal files will improve database performance. In DBMaker 3.0, previous backups cannot be used to restore a database to a specific point in time after re-sizing the journal files, however, in versions after 3.0 this is permitted. To protect a database from disk failure, perform a full backup immediately after resizing the journal files.

➲ **To resize a journal file, a DBA needs to perform the following:**

1. Estimate the disk space required to handle the largest transactions to determine the number and size of journal files required.
2. Shut down the database.
3. Update **dmconfig.ini** and re-specify these two parameters: **DB_JnFil**, **DB_JnlSz**.

   NOTE *These settings may also be changed in the advanced settings – storage page of the JServer Manager start database wizard.*
4. Set the start mode to new journal mode in **dmconfig.ini**: **DB_SMode=2**

   NOTE *This setting may also be changed in the advanced settings – start database page of the JServer Manager start database wizard.*
5. Restart the database.
6. Reset the start mode back to normal in **dmconfig.ini**: **DB_SMode=1**

   NOTE *This setting may also be changed in the start database page of the JConfiguration Tool.*
7. Perform an online full backup if a database is in *BACKUP-DATA* or *BACKUP-DATA-AND-BLOB* mode.

Tablespaces

A DBMaker database is partitioned into smaller logical areas of space known as *tablespaces*. Tablespaces are logical areas of storage that allow the database to be subdivided into manageable areas. Each tablespace contains one or more operating system files. Before starting to use tablespaces and files in DBMaker, be familiar with the terms below.

Tablespace Types

Tablespaces can be either fixed in size or automatically extensible. Tablespaces that are fixed in size are called *regular tablespaces*, and tablespaces that can have their size automatically extended are called *autoextend tablespaces*. DBMaker also has a special tablespace called the *system tablespace*.

The System Tablespace

All DBMaker databases have at least two tablespaces, one system tablespace (SYSTABLESPACE), and one default tablespace (DEFTABLESPACE). DBMaker generates a system tablespace to record the *system catalog table* whenever a database is created. The system catalog tables store information about the entire database.

The Default Tablespace

The default tablespace stores user tables when users do not specify which tablespace to be allocated. However, creating additional tablespaces for user table storage is more flexible and efficient.

Regular Tablespaces

A regular tablespace has a fixed size and contains one or more data files. If a file in a regular tablespace is too small to hold all of the data intended for it, it can be enlarged manually. The maximum number of files that can be contained in a regular tablespace is 32767. The total number of pages in all files in a tablespace must not exceed 8TB.

Autoextend Tablespaces

Autoextend tablespaces automatically grow as required. Files in an autoextend tablespace will expand automatically; DBMaker expands them by the reverse order of insertion. That means the last data file added will be the first one to expand if normal data space is required.

Any autoextend tablespace can be changed to a regular tablespace to keep the tablespace from expanding, and vice versa, a regular tablespace can be changed to an autoextend tablespace if the space is exhausted. Alternatively, new files can be added or existing files enlarged to expand a regular tablespace. Raw device files can only be used with regular tablespaces, and cannot be used with autoextend tablespaces.

DBMaker automatically creates an autoextend tablespace called the system tablespace when creating a database. When creating any other tablespaces, regular tablespaces will be used by default. To prevent the default tablespace from growing without a limit, change it to a regular tablespace.

The **dmconfig.ini** file registers the number of pages for each data file. The number of pages in a data file is the initial size of a file belonging to an autoextend tablespace, and is the actual size of a file belonging to a regular tablespace.

5.3 Managing Tablespaces and Files

There are numerous things to consider when managing tablespaces and files for a database. For example, the size and type of new tablespaces must be determined at the time of database creation, additional tablespaces can later be created, autoextend tablespaces changed to regular tablespaces and vice-versa, data files added to tablespaces, the size of files in tablespaces set and altered, and data files and tablespaces dropped when they are no longer required.

Either the JDBA Tool or a combination of dmSQL commands and modifications to the **dmconfig.ini** file can be used to manage tablespaces. The JDBA Tool provides an intuitive user interface for all tablespace management routines. For more information on how to use the JDBA Tool to manage tablespaces, refer to the *"JDBA Tool User's Guide"*.

Each DBMaker database has at least one tablespace called the system tablespace. When a database is created, DBMaker generates five files: a system data file, a user data file, a system BLOB file, a user BLOB file, and a journal file. The system data file, system BLOB file, and journal file are placed in the system tablespace. These three files are used to record the system catalog tables for the entire database. The user data file and the user BLOB file are placed in the default user tablespace.

User tables are stored in the default user tablespace unless additional tablespaces are created. Creating additional tablespaces to store user tables is more flexible and efficient.

Initial Setting of System Files and Tablespace

DBMaker generates the system tablespace and the three system files (the system data file, the system BLOB file, and the journal file) when creating a new database. These files are used to keep a record of the database schema and transactions. DBMaker concatenates the database name with the file extensions .SDB, .SBB, and .JNL to name the system data, BLOB, and journal files respectively. If the system data, BLOB, and journal file sizes are not specified, they will be created with default sizes of 800KB, 20KB, and 4000KB respectively. To use different names for the system files, specify them in the **dmconfig.ini** file, or through the storage page of the JConfiguration Tool.

➲ **Example**

To specify the names of the system files in the **dmconfig.ini** file:

```
[MY_DB]                                    ;database name
DB_DbDir = /disk1/usr                      ;database directory
DB_DbFil = datafile.sdb                    ;data file
DB_BbFil = blobfile.sbb                    ;BLOB file
DB_JnFil = jrnlfile.jnl                    ;journal file
```

If these values are in the **dmconfig.ini** file at the time the CREATE DB command is committed, then DBMaker will create the three system files as before, but this time it will use the names provided above instead of the default names. In this case, the system data file is named **datafile.sdb**, the system BLOB file is named **blobfile.sbb,** and the journal file is named **jrnlfile.jnl**.

The system tablespace is created as autoextend by default; therefore, size of the system tablespace is just an initial size, not a limitation. To limit the disk space used by the system tablespace, change the system tablespace to a regular tablespace by using the ALTER TABLESPACE command.

Once all of the space in a regular system tablespace is exhausted, the only way to enlarge it are to add files to the regular system tablespace, enlarge the system files by adding pages, or change the tablespace type to autoextend.

Initial Setting of Default User Files and Tablespace

DBMaker generates the default user tablespace and the two files (the user data file, the user BLOB file) when creating a new database. These files are used to store user data. DBMaker concatenates the database name with the file extensions .DB and .BB to name the user data and BLOB files respectively. Unless their size is specified in advance, the user data and user BLOB files will be created with default sizes of 800KB and 20KB, respectively. To use different names for the default user files, specify them in the **dmconfig.ini** file, or in the storage page of the JConfiguration Tool.

➲ **Example**

In order to specify the names of the default user files in the **dmconfig.ini** file:

```
[MY_DB]                                      ;database name
DB_UsrDb = /disk1/usr/f1.db  200             ;data file
DB_UsrBb = /disk1/usr/f1.bb  20              ;blob file
```

If a database is created with these values in the **dmconfig.ini** file, then DBMaker will create the two files using the names provided above instead of the default names. In this case, the default data file will be named **f1.db** with a size of 200 pages and the default BLOB file will be named **f1.bb** with a size of 20 frames.

The default tablespace is initially created as an autoextend tablespace, so its initial size is not a limitation.

Creating Tablespaces

Additional tablespaces can be created to contain other data and BLOB files. A tablespace may be created using dmSQL or the JDBA Tool. Details on creating tablespaces with dmSQL can be found in the *"SQL Command and Function Reference"*. Details on how to create tablespaces with JDBA Tool can be found in the *"JDBA Tool User's Guide"*.

A tablespace must contain at least one data file, but additional files in the tablespace can be either data files or BLOB files. DBMaker creates a new file as a data file by default; the file type must be specified as BLOB to create a BLOB file.

Before creating a new tablespace, specify the size and filenames of the data files associated with the tablespace in the **dmconfig.ini** file.

➲ Example 1

The following entries are required in **dmconfig.ini** to specify three files named **f1**, **f2**, and **f3** with operating system filenames and page sizes:

```
[MY DB]                           ;database name
f1 = /disk1/usr/f1.dat 1000       ;a data file with 1000 pages
f2 = /disk2/usr/f2.dat 500        ;a data file with 500 pages
f3 = /disk1/usr/f3.blb 1000       ;a blob file with 1000 pages
```

To create a regular tablespace **ts1** with two data files and one BLOB file, with the data files placed on different disks:

```
dmSQL> CREATE TABLESPACE ts1 DATAFILE f1, f2, f3 TYPE=BLOB;
```

➲ Example 2

To create an autoextend tablespace with one data file and one BLOB file. The initial size of the data file is 500 pages, and the initial size of the BLOB file is 20 pages. If the data file or BLOB file is filled, it will expand automatically:

```
[MY DB]                           ;database name
f4 = /usr/f4.dat 500              ;a data file with initial 500 pages
f5 = /usr/f5.blb 20               ;a blob file with initial 20 pages
```

To create a new tablespace that uses these files:

```
dmSQL> CREATE AUTOEXTEND TABLESPACE ts2 DATAFILE f4 TYPE=DATA, f5
TYPE=BLOB;
```

RAW DEVICE FILES

On UNIX systems, if the prefix of the physical file name is **/dev/**, DBMaker will regard it as a *raw device file*. A raw device file supports faster access than a normal file. Thus, raw device files will improve database performance. Create a raw device file on a disk before associating this file with a tablespace. Only regular tablespaces may contain raw device files.

➲ **Example**

To specify a raw device file, **f2** with the operating system filename **/dev/rawf2** with 5000 pages, add the following to **dmconfig.ini**:

```
[MY DB]                            ;database name
f2 = /dev/rawf2  5000              ;a raw device file with 5000 pages
```

To create a regular tablespace, **ts3**, containing the above raw device file:

```
dmSQL> CREATE TABLESPACE ts3 DATAFILE f2;
```

Expanding a Regular Tablespace

There are three ways to expand a regular tablespace:

- Add new files to a regular tablespace.
- Add pages to existing files in a regular tablespace.
- Set autoextend ON.

All of these functions may be performed by using the JDBA Tool or by using a combination of SQL commands and modifications to the **dmconfig.ini** file. Following is an example of how to expand a regular tablespace by editing the dmconfig.ini file and using SQL commands.

➲ **Example**

Before issuing a command, give DBMaker the name of the physical file that corresponds to the logical file named **file_blob** by adding a statement to the **dmconfig.ini** file in the section for that database. In this case, **file_blob** is the logical name that will be used in the database, and **file.blb** is the physical file name that is used by the operating system:

```
file_blob = file.blb 120
```

To add a new BLOB file named **file_blob** to a regular tablespace with 120 frames to the tablespace named **app_ts**:

```
dmSQL> ALTER TABLESPACE app_ts ADD DATAFILE file_blob TYPE = BLOB;
```

To add 100 pages to an existing data file named **file_data** in a regular tablespace named **app_ts**:

```
dmSQL> ALTER DATAFILE file_data ADD 100 PAGES
```

After altering the size of the file by adding the extra pages, DBMaker will update the number of pages for the file in the **dmconfig.ini** file to reflect the new value.

Adding Files to Tablespaces

Enlarge the size of a regular tablespace or autoextend tablespace, and consequently the database, by creating and adding new files to it. To increase the space available to insert or update data rows, add data files into a regular tablespace or autoextend tablespace. To increase the space available to store BLOB data, add BLOB files. Files may be added to a tablespace by using the JDBA Tool or by modifying the dmconfig.ini file and entering commands at the dmSQL prompt. The following is a guideline for adding files by modifying the dmconfig.ini file and entering commands at the dmSQL prompt

Be sure to first add lines to the **dmconfig.ini** file that specify the size and filenames of new files when adding data files to a tablespace. Also, specify the file type as BLOB when adding BLOB files, otherwise DBMaker will create a data file by default.

➲ **Example 1**

To specify in the **dmconfig.ini** file a data file named **f7** with 3000 pages, where the operating system filename is **/disk1/usr/f7.dat:**

```
[MY DB]                          ;database name
f7 = /disk1/usr/f7.dat 3000      ;a data file with 3000 pages
```

To add the data file **f7** into the **ts1**tablespace:

```
dmSQL> ALTER TABLESPACE ts1 ADD DATAFILE f7;
```

➲ **Example 2**

To specify in **dmconfig.ini** a BLOB file named **f8** with 5000 pages; the operating system file name is **/disk1/usr/f8.blb**:

```
[MY DB]                             ;database name
f8 = /disk1/usr/f8.blb 5000         ;a blob file with 5000 pages
```

To add this BLOB file to tablespace **ts1**:

```
dmSQL> ALTER TABLESPACE ts1 ADD DATAFILE f8 TYPE=BLOB;
```

The file type must be stated or it will be added as a data file by default.

Adding Pages to Files in Tablespaces

In addition to adding files to a regular tablespace to enlarge a database, a database can be enlarged by increasing the size of existing files in a regular tablespace. File size can be increased in autoextend tablespaces by adding pages, which pre-allocates disk space for improved performance. When the size of a file is changed, DBMaker automatically updates the entry for the file in **dmconfig.ini** to reflect the increased number of pages.

File size may be altered using the JDBA Tool or by entering the ALTER DATAFILE command at the dmSQL prompt. The following is a guideline for altering file size by entering commands at the dmSQL prompt

➲ **Example**

To alter the size and extend file **f1** by adding 100 pages, (the file **f1** must already exist and be associated with a tablespace):

```
dmSQL> ALTER DATAFILE f1 ADD 100 PAGES;
```

Changing Regular to Autoextend Tablespaces

A database administrator may want to alter a tablespace from regular to autoextend when:

- A database administrator wants to add more data to a regular tablespace, but the tablespace has already grown to fill all files belonging to this tablespace and the disk still has space.
- The database administrator does not want to restrict the amount of space a tablespace will occupy.

After creating a regular tablespace, the database administrator can change it to an autoextend tablespace by using JDBA Tool or the ALTER TABLESPACE command at the dmSQL command prompt.

➲ **Example**

To change the regular tablespace **ts1** to an autoextend tablespace:

```
dmSQL> ALTER TABLESPACE ts1 SET AUTOEXTEND ON;
```

Changing Autoextend Tablespaces to Regular Tablespaces

A database administrator may want to alter a tablespace from autoextend to regular when:

- The database administrator wants to restrict the amount of space a tablespace will occupy. An autoextend tablespace can grow to fill all available space on a disk.

After creating an autoextend tablespace, the database administrator can change it to a regular tablespace by using JDBA Tool or the ALTER TABLESPACE command at the dmSQL command prompt.

➲ **Example**

To change the autoextend tablespace, **ts1**, to a regular tablespace:

```
dmSQL> ALTER TABLESPACE ts1 SET AUTOEXTEND OFF;
```


Shrinking Tablespaces and Files

Tablespaces may be reduced in size if there is a need to allocate disk space for other uses. Two dmSQL commands can be used to reduce tablespace size, the SHRINK DATAFILE command, and the SHRINK TABLESPACE command. The SHRINK DATAFILE command works on a user-specified file, while the SHRINK TABLESPACE command works on all files in the user-specified tablespace. These operations may be carried out by using the or by using the JDBA Tool. The following sections outline how to use commands at the dmSQL command prompt to reduce tablespace size.

TRUNCATEONLY OPTION

The SHRINK command with the TRUNCATEONLY option removes contiguous free pages at the end of any data file that it is executed on. It does not compress the file; if there are free pages between used pages, they will remain in the file. The database administrator may choose to truncate all tailing free pages (without WITH *n* FREE PAGES option), or truncate free pages while still allowing a given number of free pages to remain (WITH *n* FREE PAGES option). Following are examples of both options.

Without WITH *n* FREE PAGES Option

The SHRINK command with the TRUNCATEONLY option (without WITH *n* FREE PAGES option) only truncates contiguous free pages at the end of a file.

For example, tablespace **ts1** contains **file1** and **file2**. The following diagrams, where gray blocks represent used pages and white blocks represent free pages, represent the page status of file1 and file2.

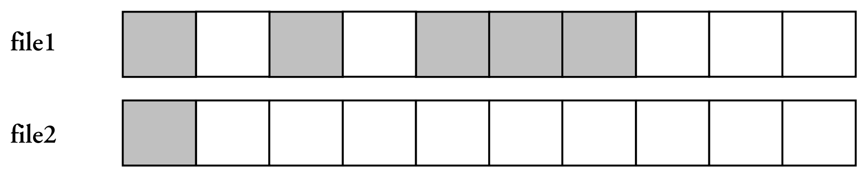

The free pages at the end of both files may be removed by executing the SHRINK TABLESPACE command on the entire tablespace, or by executing the SHRINK DATAFILE command on both files. The TRUNCATEONLY option must be specified. The following examples demonstrate.

- **Example 1**

```
dmSQL> SHRINK TABLESPACE ts1 TRUNCATEONLY;
```

- **Example 2**

```
dmSQL> SHRINK DATAFILE file1 TRUNCATEONLY;
dmSQL> SHRINK DATAFILE file2 TRUNCATEONLY;
```

After truncating, the pages at the end of both files have been removed. A graphical representation of the page status of both files follows:

Result

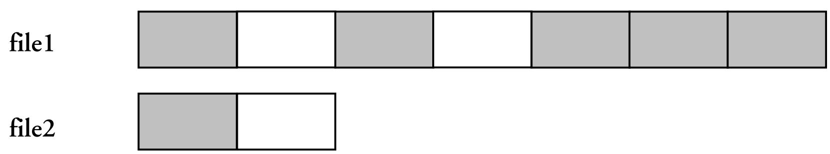

Although all pages of **file2** are free, DBMaker reserves at least two pages (one is a PE page and one is a data page).

WITH *n* FREE PAGES Option

The WITH *n* FREE PAGES option specifies the total number of tailing free pages (not including the PE page) to remain in the file after it has been truncated.

Using the previous example of **file1** and **file2**, execute one of the following.

- **Example 1**

```
dmSQL> SHRINK TABLESPACE ts1 TRUNCATEONLY WITH 3 FREE PAGES;
```

- **Example 2**

```
dmSQL> SHRINK DATAFILE file1 TRUNCATEONLY WITH 3 FREE PAGES;
dmSQL> SHRINK DATAFILE file2 TRUNCATEONLY WITH 3 FREE PAGES;
```

Result:

file1

T

file2

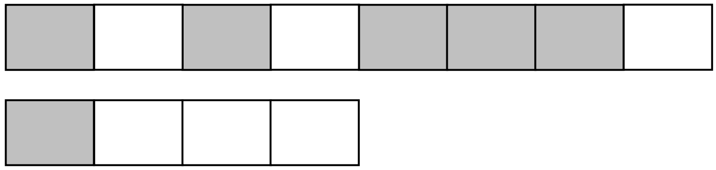

e

SHRINK TABLESPACE command and the WITH FREE PAGES option apply individually to each file in a tablespace. In the above case, there are three free pages reserved for each file in the same tablespace.

It is not possible to inadvertently add pages to a file by specifying more free tailing pages than the file currently has. For example, if there are 50 free pages in a file, specifying the option WITH 80 FREE PAGES option causes nothing to happen. After the SHRINK command, there are still 50 free pages and it does not enlarge the file size by adding 30 (80 - 50) free pages.

The SHRINK command should be executed with autocommit ON. The TRUNCATEONLY option cannot be rolled back. Users cannot roll back this command, even through crash-recovery.

COMPRESSONLY OPTION

Only the SHRINK TABLESPACE command supports the COMPRESSONLY option. It compresses each file in the tablespace. It does not compress records on the same page because the smallest unit used for compression is a page. It moves the used pages in tail of the file to free front pages. After using the command, all free pages are placed at the end of the file and all used pages at the front.

Result 1:

file1

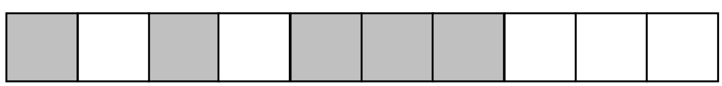

File1 has four used pages that are not adjacent.

➲ **Example**

To make it contiguous:

```
dmSQL> SHRINK TABLESPACE ts1 COMPRESSONLY;
```

Result 2:

file1

The SHRINK command must be executed with autocommit ON. The COMPRESSONLY option can be rolled back. If the database crashes, the operation of COMPRESSONLY will be all done or all failure after crash-recovery.

There are some conflicts between the SHRINK command with the COMPRESSONLY option and using backup. DBMaker does not allow these two commands to be executed at the same time.

Limitations to Shrinking and Compressing Tablespaces

The general limitations for these commands are:

- The SHRINK command can be used on data and BLOB files but not on a journal file.
- Only a user with DBA authority can execute the SHRINK command.
- The SHRINK command requires autocommit ON.
- The SHRINK command was added in DBMaker 3.7; early versions of DBMaker do not recognize this command. Therefore, once a DBA executes the SHRINK command and does the incremental backup with it, the earlier version of DBMaker cannot restore the journal backup file.
- The TRUNCATEONLY option cannot be rolled back.
- The COMPRESSONLY option cannot compress the SYSTABLESPACE tablespace.
- The COMPRESSONLY option does not check if user tables have an OID column or not. An OID column is used to reference a record elsewhere in the

database. After using COMPRESSONLY, an OID column may no longer point to the correct record if the referenced record is in the compressed tablespace or file. It does not modify OID columns in user tables.

- The COMPRESSONLY option and backup command cannot be executed at the same time.

Dropping Tablespaces

If a tablespace is empty or contains information that is no longer required, a database administrator can drop it from the database. Any tablespace in a DBMaker database, except the system tablespace, can be dropped. To drop a tablespace, first drop all tables in the tablespace or ensure it is already empty of tables. For more information on how to drop tables from a tablespace, refer to Chapter 6, *Managing Schema Objects*.

Dropping a tablespace will automatically drop all the files associated with it, but will not remove them from the file system of the operating system. Those files will still exist in the file system and can only be removed using operating system commands to recover the disk space they occupy. The data stored in the physical files corresponding to a tablespace is not recoverable once the physical files have been removed from the file system. Be careful when removing files associated with tablespaces or valuable data may be lost.

Tablespaces may be dropped using JDBA Tool or by using the DROP TABLESPACE command at the dmSQL command prompt.

➲ **Example**

To drop the tablespace **ts2** and all files associated with it:

```
dmSQL> DROP TABLESPACE ts2;
```

Dropping Files From a Tablespace

Users are able to remove unwanted datafiles from a tablespace by using JBDA Tool or by using the ALTER TABLESPACE *tablespace-name* DROP DATAFILE *file-name* command at the dmSQL command prompt. Unwanted datafiles can be removed from a tablespace with the following conditions:

a) Users are not able to drop a datafile from a tablespace if it is the only datafile in that table space.

b) The datafile to removed must be empty.

c) Users cannot remove the system or default datafile from the system or default tablespace.

➲ **Example**

To drop the datafile **df2** from a tablespace **ts2**:

```
dmSQL> ALTER TABLESPACE ts2 DROP DATAFILE df2
```

Getting Information about Tablespaces and Files

Using JDBA Tool, it is straightforward to view the structure of tablespaces and files within a given tablespace. Tablespaces are displayed as part of the logical tree structure of all database objects. Selecting the tablespaces node on the tree will expand the tree to display all tablespaces in the database. Selecting a tablespace from the tree will display all files in the tablespace as well as details about the files, such as size, physical location, data type, or whether the tablespace is extensible.

Alternatively, use dmSQL to select all columns of the system table SYSTABLESPACE for information on tablespaces, or SYSFILE for information on user BLOB and data files.

➲ **Example 1**

To obtain information on tablespaces, such as tablespace names, whether they are regular or autoextend tablespaces, the number of files associated with tablespaces, and the number of total pages, browse the system table SYSTABLESPACE in the system catalog:

```
dmSQL> SELECT * from SYSTABLESPACE;
```

Example 2

To obtain information about files in a similar manner by browsing the system table SYSFILE to get information about file names, file types, database internal file identification, which tablespace files are associated with and how many pages each file contains:

```
dmSQL> SELECT * from SYSFILE;
```

For more information about the system catalog tables SYSTABLESPACE and SYSFILE, refer to "*Appendix B*".

Checking File and Tablespace Consistency

DBMaker supports six commands to check the consistency of different parts of a database. These commands are time consuming when the database is large and they will take locks, and should only be used when necessary. File and tablespace consistency may be checked using one of these commands. The CHECK FILE command will check if a file is corrupted or if a tablespace contains the correct tables.

CHECKING FILES

DBMaker allows the contents of every page or frame in a data file to be checked. Any corruption found when checking files is usually caused by disk errors.

Example

To check consistency for the **FILE1** data file:

```
dmSQL> CHECK FILE FILE1;
```

CHECKING TABLESPACES

DBMaker allows files and tables associated with a tablespace to be checked. When checking files and tables, DBMaker uses the same methods as the check file and check table commands, and returns the same results as if these commands were executed directly.

➲ **Example**

To check tablespace consistency for the **TS1** tablespace:

```
dmSQL> CHECK TABLESPACE TS1;
```

6 Managing Schema and Schema Objects

This chapter discusses the management of different types of schema objects in DBMaker, including tables, views, synonyms, indexes, serial numbers, data integrity, and domains.

The chapter includes topics on browsing the system catalogs to get information about schema objects, and how to estimate the disk storage space required for tables and indexes.

Schema object management may be carried out by using dmSQL commands or through the JDBA Tool. The JDBA Tool contains an intuitive graphical interface, provides easy-to-use wizards for most database management tasks, and displays the logical structure of the database in an unambiguous format. Using JDBA Tool will aid first time users of DBMaker in understanding the relationship between schema objects. Experienced users will find the logical display aids in the creation and management of database schema. The following sections show examples of how to manage database schema objects though dmSQL. For more information on using JDBA Tool to manage schema objects, refer to the "*JDBA Tool User's Guide*".

6.1 Managing Schema

Schema are namespaces (logical grouping of database objects). Schema contain schema objects such as tables, views, indexes, commands, procedures and a domain and a synonym.

CREATE SCHEMA defines a new schema. After the schema is created, we can create objects within the schema. The schema owner is the grantor for any privileges granted.

The owner of the schema is determined as follows:

- If an AUTHORIZATION clause is specified, the specified user-name is the schema owner. And if schema-name is omitted, the specified user-name is used as the schema name.

For example

```
CREATE SCHEMA AUTHORIZATION JEFFERY.
```

- If an AUTHORIZATION clause is not specified, the user that issued the CREATE SCHEMA statement is the schema owner.

➲ **Example 1**

As a user with RESOURCE authority, JEFFERY, creates a schema called **SS1**. JEFFERY is the default owner.

```
CREATE SCHEMA SS1
```

➲ **Example 2:**

As a user with DBA authority, creates a schema with the user JEFFERY as the owner, and the username JEFFERY is the default schema name.

```
CREATE SCHEMA AUTHORIZATION JEFFERY
```

➲ **Example 3:**

As a user with DBA authority, creates a schema called **SS2** with the user JEFFERY as the owner.

```
CREATE SCHEMA SS2 AUTHORIZATION JEFFERY
```

➲ **Example 4:**

A user with DBA authority creates a schema, **inventory**. The user then creates a table and an index on that table. The user final grants authority on the table to the user JEFFERY.

```
CREATE SCHEMA inventory;
CREATE TABLE inventory.part (partNo smallint not null, quantity int);
CREATE INDEX partind ON inventory.part (partNo);
GRANT ALL ON inventory.part TO JEFFERY;
```

DROP SCHEMA removes schemas from the database. A schema can only be dropped by its owner or a DBA. Note that the owner can't drop the schema if it contains any objects.

➲ **Example:**

Remove schema **oldclients** from the database.

```
DROP SCHEMA oldclients
```

6.2 Managing Tables

Tables are the logical unit of storage used by DBMaker to store data. A table consists of several columns and rows. A column is sometimes referred to as a field or attribute, and a row can be referred to as a record or tuple.

In DBMaker, each table is identified by a unique schema name and table name.

For example, if two users called **Jeff** and **Kevin** each create a table named **friend** with the default schema name, then the table names **Jeff.friend** and **Kevin.friend** denote the two different tables.

In the JDBA Tool, all tables in a database can be viewed by expanding the tables node on the logical tree. Selecting a table displays that table's schema.

Creating Tables

Every table is defined with a table name and a set of columns. The number of columns in a table can range from 1 to 252.

Each column has:

- A column name and a data type (or a domain, which is described in section 6.10, Managing Domains)
- A length (the length might be predetermined by the data type, such as INTEGER), a precision and scale (for columns of the DECIMAL data type only) or a starting number (for columns of SERIAL data type only)

DBMaker supports a large number of data types that can be used to define columns. There are numerical types (SMALLINT, INTEGER, FLOAT, DOUBLE, DECIMAL and SERIAL), binary types (BINARY, VARBINARY, CHAR, and VARCHAR), BLOB types (LONG VARCHAR, LONG VARBINARY and FILE), and time types (DATE, TIME and TIMESTAMP). See the "*SQL Command and Function Reference*" for more information about data types.

When creating a table, provide the table name, column definitions, and the name of the associated tablespace. A table will be placed in the system tablespace by default if it is not associated with another tablespace. Tables may be created using the JDBA Tool Create Table wizard or using the dmSQL command prompt. For information on creating a table with JDBA Tool, refer to the "*JDBA Tool User's Guide*". The following is an example of how to create a table using dmSQL. Details on syntax and usage of the SQL command CREATE TABLE can be found in the "*SQL Command and Function Reference*".

➲ **Example**

To create the **employee** table in tablespace **ts1**:

```
dmSQL> CREATE TABLE employee (nation CHAR(20),
                                  ID INTEGER,
                                name CHAR(30),
                            joinDate DATE,
```

```
                          height FLOAT,
                          degree VARCHAR(200),
                         picture LONG VARCHAR) IN ts1;
```

DBMaker provides many useful features that can be applied when creating tables:

- Defining a default value for a column.
- Specifying that a column is not nullable.
- Specifying the primary key or the foreign key for the table.
- Specifying the LOCK MODE, FILLFACTOR, or NOCACHE options to improve database efficiency.
- Specifying the table as temporary.
- Specifying the table to automatically update statistics.

DEFAULT VALUES FOR COLUMNS

A column in a table can be assigned a default value so that when a new row is inserted and a value for the column is omitted, the default value will be automatically supplied.

Default values for each column in a table may be specified. If a default value is not defined for a column, the default value for the column is set to NULL.

Legal default values can be constants or built-in functions. For more information about built-in functions, refer to the *"SQL Command and Function Reference*". The following example shows how to specify the default value of a column as a built in function using dmSQL.

➲ **Example**

To specify the default value of the column **nation,** in the table **employee** as a constant—'R.O.C.' and the default value of the column **joinDate** as the value of the built-in function **curdate**():

```
dmSQL> CREATE TABLE employee (nation CHAR(20) DEFAULT 'R.O.C',
                               ID INTEGER,
                             name CHAR(30),
                         joinDate DATE DEFAULT CURDATE(),
```

```
                               height FLOAT,
                               degree VARCHAR(200),
                              picture LONG VARCHAR) IN ts1;
```

NOT NULL

Rules for columns or tables may be specified. These rules are called *integrity constraints*. One example is the NOT NULL integrity constraint defined on a column in a table. It enforces the rule that the column cannot contain a null value.

For example, the **employee** table might always need an ID and a name for a new employee.

➲ **Example**

To create an ID and name for new employees on the **employee** table:

```
dmSQL> CREATE TABLE employee (nation CHAR(20) DEFAULT 'R.O.C',
                                   ID INTEGER NOT NULL,
                                 name CHAR(30) NOT NULL,
                             joinDate DATE DEFAULT CURDATE(),
                               height FLOAT,
                               degree VARCHAR(200)) IN ts1;
```

PRIMARY KEY AND FOREIGN KEYS

The table owner can specify the primary key or foreign key with the CREATE TABLE command. Refer to section 6.8, *Managing Data Integrity* for information on primary and foreign keys.

LOCK MODE

The *lock mode* of a table identifies the type of lock that DBMaker automatically places on objects when accessing the database. DBMaker supports three lock mode levels: TABLE, PAGE, and ROW. The PAGE lock mode is used by default if the lock mode is not specified when a table is created. If the lock mode is set to a higher level (such as TABLE), the level of concurrency on database accesses will be lower, but the required lock resources (shared memory) will also be smaller. If the lock mode is set to a lower level (such as ROW), the level of concurrency on database accesses will be higher, but the required lock resources (shared memory) will be larger. In other words, if a user inserts or modifies rows in a table with the lock mode set to TABLE, no one else will be able to access the table. The reason for this is that an exclusive lock is taken on the entire table. For more information about lock modes, see section 9.3, "Locks".

➲ **Example**

To specify the lock mode on a table:

```
dmSQL> CREATE TABLE employee (nation CHAR(20) DEFAULT 'R.O.C',
                                  ID INTEGER NOT NULL,
                                name CHAR(30) NOT NULL,
                            joinDate DATE DEFAULT CURDATE(),
                              height FLOAT,
                              degree VARCHAR(200)) IN ts1
                    LOCK MODE ROW;
```

FILLFACTOR

The FILLFACTOR feature optimizes the utilization of space for data pages by reserving space for the expansion of existing records. It specifies the percentage of a page that can be filled before stopping new records from being inserted. Using this method records can be accessed more efficiently by avoiding the need to retrieve information for one record from multiple pages.

➲ **Example**

To set the FILLFACTOR of the **employee** table to be 80%:

```
dmSQL> CREATE TABLE employee (nation CHAR(20) DEFAULT 'R.O.C',
                                  ID INTEGER NOT NULL,
                                name CHAR(30) NOT NULL,
                            joinDate DATE DEFAULT CURDATE(),
                              height FLOAT,
                              degree VARCHAR(200)) IN ts1
                    LOCK MODE ROW
                    FILLFACTOR 80;
```

In this case, new rows cannot be inserted into the data page after the used space is larger than 80%. The legal values for the FILLFACTOR can be from 50 to 100, and the default value is 100.

NOCACHE

The NOCACHE feature is useful when accessing large tables with a table scan. Although DBMaker uses page buffers in shared memory to cache retrieved data and avoid frequent disk I/O, table scans on large tables can still cause frequent disk I/O activity. This happens during a table scan on a table with a larger number of data pages than the number of page buffers, which causes all page buffers to be exhausted.

Once the NOCACHE option is specified when creating a table, DBMaker only uses one page buffer to cache the data retrieved from a table during a table scan. This prevents the page buffers from being exhausted by only one large table scan.

➲ **Example**

To specify the NOCACHE option:

```
dmSQL> CREATE TABLE employee (nation CHAR(20) DEFAULT 'R.O.C',
                                  ID INTEGER NOT NULL,
                                name CHAR(30) NOT NULL,
                            joinDate DATE DEFAULT CURDATE(),
                              height FLOAT,
                              degree VARCHAR(200)) IN ts1
                    LOCK MODE ROW
                    FILLFACTOR 80
                    NOCACHE;
```

TEMPORARY TABLES

A *temporary table* may be created for storing data. Temporary tables only exist during a single session, and DBMaker will automatically drop the temporary table when the user that created it is disconnected from the database. Temporary tables support fast data operations and can be used only by the creator. Operations that are executed improperly on a temporary table will result in incorrect data on the table. Client users may also create a local temporary table using the CREATE LOCAL TEMPORARY TABLE syntax.

Example 1

To create a temporary table named **student**:

```
dmSQL> CREATE TEMPORARY TABLE student (name CHAR(25) NOT NULL,
                                       birthday DATE,
                                       score INTEGER);
```

Example 2

To create a local temporary table named **student**:

```
dmSQL> CREATE LOCAL TEMPORARY TABLE student (name CHAR(25) NOT NULL,
                                             birthday DATE,
                                             score INTEGER);
```

AUTO UPDATE STATISTICS

The statistics value is very important if a table is read frequently. DBMaker can automatically update the statistics of a table periodically if a user has specified the update statistics time interval when creating it.

Example

To create a table and specify to automatically update its statistics every 7 days:

```
dmSQL> CREATE TABLE student (name CHAR(25) NOT NULL,
                                       birthday DATE,
                                       score INTEGER)
       SET UPDATE STATISTICS EVERY 7 DAYS;
```

Statistics are only updated after the database has been started. Updating statistics also requires processor resources and will affect database performance. Selecting an interval and a time that does not interfere with peak table usage will prevent degradation of performance while still providing updated statistics.

Browsing Table Schema

The schema of a table may be queried by using dmSQL or the JDBA Tool. JDBA Tool provides a graphical representation of table schema and allows table schema to be modified without entering any SQL commands. It is also possible to use the dmSQL command DEF TABLE to directly query a table's schema.

➲ **Example**

To view the schema for table **supportqueries**:

```
dmSQL> DEF TABLE SUPPORTQUERIES;

create table SYSADM.SUPPORTQUERIES (

 LOGINID  CHAR(20) default null ,

 REQUEST  LONG VARCHAR default null ,

 REQUESTTIME  TIMESTAMP default null ,

 ATTACHMENT  LONG VARBINARY default null )

 in DEFTABLESPACE  lock mode page  fillfactor 100 ;

create  text index REQUEST on SYSADM.SUPPORTQUERIES ( REQUEST )  text
block size 500  basic bit length 1024  extended bit length 1024  cluster
width 4096  ;\
```

Altering Tables

After a table is created in DBMaker, a user with modify permission can alter it by:

- Adding/dropping columns
- Modifying column definitions
- Changing the FILLFACTOR value
- Turning on/off the NOCACHE option

The schema of a table may be altered using dmSQL commands or with JDBA Tool.

Adding/Dropping Columns

A user with modify permission can add/drop one or multiple columns in a table whether the column is empty or not. Adding a new column to an empty table is the same as expanding the table schema and placing the new column in the last position. A user with modify permission can also add a new column before or after any existing column in a table.

When adding a new column to a table, DBMaker not only expands the table schema but also fills all rows in the new column with NULL values by default. If a user with modify permission wants to add a column with the NOT NULL integrity constraint to a table, give a specified value for the existing records on the column (a default value, as mentioned in "Default Values for Columns"). For detailed SQL syntax, refer to the "*SQL Command and Function Reference*".

➲ **Example 1**

To add a column named **photo** to the **employee** table:

```
dmSQL> ALTER TABLE employee ADD COLUMN photo LONG VARCHAR;
```

➲ **Example 2**

To add a column named **city** after the existing column name to the employee table and set the default value to 'Taipei':

```
dmSQL> ALTER TABLE employee ADD COLUMN city CHAR(20) default 'Taipei'
AFTER name;
```

➲ **Example 3**

If the **employee** table is not empty and a user wants to add a non-null column to it, the GIVE keyword can be used to specify a value for the existing records on the newly added column. To add a non-null column named **HireDate** to the **employee** table:

```
dmSQL> ALTER TABLE employee ADD COLUMN (HireDate NOT NULL give '2000-02-
20';
```

➲ **Example 4**

To drop a column named **photo** from the **employee** table:

```
dmSQL> ALTER TABLE employee DROP COLUMN photo;
```

MODIFYING COLUMN DEFINITION

The definition of every existing column in a table can be altered, such as column name, data type, column order, default value, column constraint, etc. Before modifying the data type of one column, make sure that the new data type is compatible with the original one, or the modifying operation will fail due to data incapability. For example, a CHAR type data column cannot be modified to a DATE type data column.

- **Example 1**

To modify the column named **photo** in the **employee** table:

```
dmSQL> ALTER TABLE employee MODIFY photo NAME TO emp_photo;
```

- **Example 2**

To modify the data type for a column named **height** in the **employee** table:

```
dmSQL> ALTER TABLE employee MODIFY height TYPE TO decimal(10,2);
```

- **Example 3**

To modify the column order for a column named **height**, place it before the **HireDate** column:

```
dmSQL> ALTER TABLE employee MODIFY height BEFORE HireDate;
```

- **Example 4**

To modify the default value for a column named **nation**:

```
dmSQL> ALTER TABLE employee MODIFY nation DEFAULT TO 'Taiwan';
```

- **Example 5**

To modify the constraint for a column named **height**:

```
dmSQL> ALTER TABLE employee MODIFY height CONSTRAINT TO CHECK value <
250;
```

Changing the Lock Mode

To gain a higher level of concurrency on simultaneous connections to a database, set the lock mode to a lower level (such as a **ROW** lock). However, doing this causes DBMaker to expend more resources; deciding which lock mode to use on a table always involves a trade-off. For more information about lock modes, see section 9.3, "Locks".

➲ **Example**

To change the lock mode for the **employee** table:

```
dmSQL> ALTER TABLE employee SET LOCK MODE ROW;
```

CHANGING THE FILLFACTOR VALUE

FILLFACTOR may be specified during table creation or later modified. For more information on the FILLFACTOR option, refer to the subsection "FILLFACTOR" in "Creating Tables".

➲ **Example**

To change the FILLFACTOR value for a table:

```
dmSQL> ALTER TABLE employee SET FILLFACTOR 90;
```

TURNING NOCACHE ON/OFF

The ON/OFF option can be used at any time for NOCACHE. For more information on the NOCACHE option refer to the subsection *NOCACHE* in.

➲ **Example**

To turn the NOCACHE option for a the **employee** table OFF:

```
dmSQL> ALTER TABLE employee SET NOCACHE OFF;
```

Locking Tables

Although DBMaker automatically handles the lock mechanism whenever a database is accessed, a table may be manually locked for subsequent SELECT or UPDATE statements. Locking a table while a user is viewing or modifying it will prevent updates by other people.

DBMaker supports some options for locking tables, such as *shared locks* for viewing data or *exclusive locks* for modifying data, and the WAIT or NO WAIT option which is used when obtaining a lock. For more information about these features, see the "*SQL Command and Function Reference*". To learn about table locks, concurrency control, and transaction handling, refer to Chapter 9, *Concurrency Control.*

➲ **Example**

To lock the **employee** table for later selections and not wait if it cannot get the table lock right away:

```
dmSQL> LOCK TABLE employee IN SHARE MODE NO WAIT;
```

Dropping Tables

A user can drop a table when the table is not being used any more. When a table is dropped, all data and indexes for this table are dropped, and pages allocated by the dropped table are released.

➲ **Example**

To drop the **employee** table:

```
dmSQL> DROP TABLE employee;
```

6.3 Managing Views

DBMaker provides an ability to define a *virtual table* called a *view*, which is based on existing tables and is stored as a definition with a user-defined view name. The view definition is stored persistently in the database, but the actual data is not physically stored there. Rather, the data is stored in the original base tables the views were derived from. A view is defined by a query that references one or more tables or other views.

Views are a very helpful mechanism in a database. For example, complex queries can be defined once and used repeatedly without having to be re-written. Furthermore, views can be used to enhance the security of a database by restricting access to a predetermined set of rows and/or columns in a table.

A user cannot determine from a view which rows of tables to update, since a view is derived from queries on tables. Due to this limitation, views can only be queried. Users can update views, but cannot insert into or delete from views.

Creating Views

Views may be created with dmSQL or the JDBA Tool. A view is defined by a name together with a query that references tables or other views. The query that defines a view cannot contain the ORDER BY clause or UNION operator.

A user can specify a list of column names for a view. If the user does not specify the column names, the view will inherit the column names of the underlying tables.

For example, to allow other users to see only two columns from the **employee** table, create a view with the SQL command shown below. Users can then view only two columns, (**name** and **ID**), from the **employee** table through the view **empView**.

➲ **Example**

Use CREATE VIEW to create the view **empView** from the employee table:

```
dmSQL> CREATE VIEW empView (empName, empId) AS
              SELECT name, ID FROM employee;
```

Browsing View Schema

The construction of a view may be queried by using dmSQL or the JDBA Tool. Use the dmSQL command DEF VIEW to directly query a table's schema.

➲ **Example**

To view the construction for view **USER_DATA**:

```
dmSQL> DEF VIEW USER DATA; create view SYSADM.USER DATA as select
USERDATA.FIRSTNAME, USERDATA.LASTNAME, USERDATA.ADDRESS1,
USERDATA.CITY1, USERDATA.COUNTRY1, USERSTATUS.USERSTATUS,
SUPPORTQUERIES.REQUESTTIME from SYSADM.USERDATA, SYSADM.USERSTATUS,
SYSADM.SUPPORTQUERIES ;
```

Dropping Views

A view can be dropped when it is no longer required. When a view is dropped, only the definition stored in the system catalog is removed. The base tables that the view was derived from are unaffected.

➲ **Example**

To drop the view **empView** use the DROP VIEW command:

```
dmSQL> DROP VIEW empView;
```

6.4 Managing Synonyms

A *synonym* is an alias for any table or view. Since a synonym is simply an alias, it requires no storage other than a definition in the system catalog.

Synonyms are useful for simplifying a fully qualified table or view name. DBMaker normally identifies tables and views with fully qualified names that are composites of the owner and object names. By using a synonym, anyone can access a table or view using the corresponding synonym without having to make use of the fully qualified name. Because a synonym has no owner name, each synonym in the database must be unique so DBMaker can identify them. Synonyms may be created or dropped with dmSQL or the JDBA Tool

Creating Synonyms

➲ **Example**

To create a synonym **employee** for the table **TOM.employee**:

```
dmSQL> CREATE SYNONYM employee FOR TOM.employee;
```

Assume that the owner of the table **employee** is **TOM**. This command creates the alias **employee** for the table **TOM.employee**. All database users can directly reference the table **TOM.employee** through the synonym **employee**.

Dropping Synonyms

A synonym that is no longer required can be dropped. When a synonym is dropped, only its definition is removed from the system catalog.

➲ **Example**

To drop the **employee** synonym:

```
dmSQL> DROP SYNONYM employee;
```

6.5 Managing Indexes

An index provides support for fast random access to a row. Building indexes for a table speeds up searching. For example, when a user executes the query SELECT NAME FROM EMPLOYEE WHERE id=306004, it is possible to retrieve the data in a much shorter time if there is an index created for the **ID** column.

An index can be composed of more than one column, up to a maximum of 16. Although a table can have up to 252 columns, only the first 127 can be used in an index.

An index can be *unique* or *non-unique*. In a unique index, no more than one row can have the same key value, with the exception that any number of rows may have NULL values. If a user creates a unique index on a table, DBMaker will check whether all existing keys are distinct or not. If there are duplicate keys, DBMaker will return an error message. After creating a unique index on a table, if a user inserts a row in the table and DBMaker will ensure that there are no existing rows with the same key as the new row.

When creating an index, the sort order of each index column can be specified as ascending or descending. For example, suppose there are five keys in a table with the values 1, 3, 9, 2, and 6. In ascending order, the sequence of keys in the index is 1, 2, 3, 6, and 9, and in descending order, the sequence of keys in the index is 9, 6, 3, 2, and 1.

When a user implements a query, the index order will occasionally affect the order of the data output.

➲ **Example**

If the following query is executed:

```
dmSQL>select name, age from friend_table where age > 20
```


Using an index with a descending order on the column **age**, the output will appear as follows:

```
      name                      age
----------------          ----------------
Jeff                                    49
Kevin                                   40
Jerry                                   38
Hughes                                  30
Cathy                                   22
```

A user can specify the *fill factor* for tables when creating an index. The fill factor denotes how dense the keys will be in the index pages. The legal fill factor values are in the range from 1% to 100%, and the default is 100%. If a user updates data often after creating the index, the user can set a loose fill factor in the index, for example 60%. If the user never updates the data in this table, then the fill factor can be left at the default value of 100%.

A user may also specify to create an index on a separate tablespace. This can result in improved disk I/O for searches that use the index if multiple disks are used.

Before creating indexes on a table, it is recommended to load all data first, especially if there is a large amount of data for that table. If a user creates an index before loading the data into a table, the indexes will be updated each time the user loads a new row. It is far more efficient to create an index after loading a large amount of data than to create an index before loading the data.

Creating Indexes

Indexes may be created using the Create Index wizard of the JDBA Tool, or the dmSQL CREATE INDEX command. To create an index on a table, specify the index name and index columns. Specify the sort order of each column as ascending or descending. The default sort order is ascending.

➲ **Example 1**

To create an index, **idx1***,* on the column **ID** for the table **salary** in descending order use the DESC option:

```
dmSQL> CREATE INDEX idx1 ON salary (ID DESC);
```

➲ **Example 2**

To create a unique index, idx2*,* on the column **ID** for the table **salary** use the UNIQUE option:

```
dmSQL> CREATE UNIQUE INDEX idx2 ON salary (ID);
```

➲ **Example 3**

To create an index with a specified fill factor use the FILLFACTOR option:

```
dmSQL> CREATE INDEX idx3 ON salary(name, age DESC) FILLFACTOR 60;
```

➲ **Example 4**

To create an index on tablespace **ts1**:

```
dmSQL> CREATE INDEX idx4 ON salary(name, age DESC) IN ts1 FILLFACTOR 60;
```

Dropping Indexes

Indexes may be dropped using the JDBA Tool, or the dmSQL DROP INDEX statement. If the index is a primary key and is referred to by other tables it cannot be dropped. For information on primary keys, refer to the section *Managing Data Integrity*.

➲ **Example**

To drop the index **idx1** from the table **salary**:

```
dmSQL> DROP INDEX idx1 FROM salary;
```

Rebuilding Indexes

Indexes may be rebuilt using JDBA Tool, or the dmSQL REBUILD INDEX statement. In general, the index will need to be rebuilt if it becomes fragmented, which reduces its efficiency. Rebuilding an index will drop the old index and then create a new one.

➲ **Example**

To rebuild the index **idx1** for the table **salary**:

```
dmSQL> REBUILD INDEX idx1 FOR salary;
```

6.6 Managing Text Indexes

A *text index* is a mechanism that provides fast access to rows in a table that contains one or more words or phrases in columns. Text indexes contain a representation of all the text found in the columns they are based on, but the data is encoded and structured to make retrieval much faster than directly from the table. Once a user creates a text index for a table, its operation is transparent. The DBMS uses the index to improve full-text query performance whenever possible.

DBMaker provides two text indexing methods: signature and inverted file (IVF).

Text indexes can be built on all character type columns, including CHAR, VARCHAR, LONG VARCHAR, LONG VARBINARY, and FILE data types. A table can have many text indexes and a text index can be built using multiple columns. A user may create text indexes by using either the JDBA Tool or the CREATE [SIGNATURE | IVF] TEXT INDEX dmSQL command.

➲ **Example**

To use a text index in the **data** column automatically (without specifying it):

```
dmSQL> SELECT id FROM book WHERE data MATCH 'compute';
```

The string operators for DBMaker include MATCH, CONTAIN, CONTAINS, and LIKE. Only the MATCH and CONTAINS operators can be applied to a text index search.

DBMaker provides two different types of text index: signature and inverted-file. Signature text index is more efficient for small amount of data. Inverted-file text index usually comsumes more storage space but its response for queries is faster for large amount of data

Creating Signature Text Indexes

DBMaker creates signature text indexes if no text index method is specified in the command. A user may create text indexes by using either the JDBA Tool, the CREATE TEXT INDEX dmSQL command, or the CREATE SIGNATURE TEXT INDEX command.

➲ **Example**

To create a signature text index, **ix_data**, on the data column for the table **book**:

```
dmSQL> CREATE SIGNATURE TEXT INDEX ix_data ON book(data);
```

SIGNATURE TEXT INDEX PARAMETERS

DBMaker provides two parameters for conveniently configuring performance and storage size of signature text indexes

- **Total text size (MB)** — the estimated total size of all source documents,in megabytes (MB). The range is 1 ~ 200 and the default is 32. Please note the real total text size is not limited to 200 megabytes; if the size is larger than 200, set to 200. However, we strongly recommend using IVF text index to index very large amount of data for significantly better query performance.
- **Scale** — the expected index size-to-total text size ratio. If a user sets **total text size** to 20(MB) and expects the text index to use 10MB of storage, then he should set **scale** to 50 (50%). The larger **scale**, the better search performance. The range is 10 ~ 200 and the default value is 40 (40%)

➲ **Example**

To create a text index, **ix_data** on the column **data** of the table **book** that contains about 40 megabytes of data, and we wish the text index uses about 20 megabytes of storage space:

```
dmSQL> CREATE SIGNATURE TEXT INDEX ix_data ON book(data)
```



```
    TOTAL TEXT SIZE 40 MB
    SCALE 50;;
```

A user can use the default setting as text index parameters. To get higher text index performance or to reduce the text index size, change the text index parameters. Set the parameters and monitor the text index performance, and then re-adjust the parameters.

Creating Inverted-File (IVF) Text Indexes

A user can create inverted-file text indexes by using the CREATE IVF TEXT INDEX command.

➲ **Example**

To create an inverted-file text index, **ix_data**, on the data column for the table **book**:

```
dmSQL> CREATE IVF TEXT INDEX ix_data ON book(data);
```

Inverted-File Text Index Parameters

There are two parameters for use in the creating IVF text index command:

Storage path — the logical working directory where the inverted-files will reside in. Users should define the logical directory in the dmconfig.ini file. The default is the value of DB_DbDir, the database's home directory. The detail storage management and naming convention of inverted-file index will be described in the next section.

Total text size (**MB**) — the approximate total size of documents will be indexed in the future. The unit of size is mega-byte (MB). Based on the size, DBMaker will decide how many partitions will be made. It may range between 1 MB to 10000 MB , and the default value is 500 MB.

➲ **Example**

To create an inverted-file text index, **ix_data** in the path \IVFDIR on the column **data** of the table **book** that contains about 400 megabytes of data:

First, add a logical path in the database's dmconfig.ini section.

```
MYPATH1 = \IVFDIR
```

Use the following command.

```
dmSQL> CREATE IVF TEXT INDEX ix data ON book(data)
    2> STORAGE PATH MYPATH1
    3> TOTAL TEXT SIZE 400 MB;
```

While creating inverted-file text indexes, it requires large amount of memory resource. DBMaker will take a simple rule to decide the maximum memory usage for creating text indexes. If DBMaker cannot detect free memory or free memory resource less than 128MB, then the maximum memory usage will be 64MB, otherwise will be half of free memory resource. Users can specify the approximate upper bound of memory usage manually through dmconfig.ini by adding a keyword entry DB_IFMem in mega-byte (MB).

➲ **Example**

To specify 100MB memory usage for creating inverted-file text-index in dmconfing.ini:

```
DB_IFMEM = 100
```

STORAGE OVERVIEW

In addition to the working directory specified by the **Storage path** parameter, DBMaker will generate sub-directories in this directory to manage different inverted-file indexes. Each inverted-file index has a unique time version, so DBMaker can use this property to generate a unique sub-directory to store index files. How to name the sub-directory is described later. This is also a limitation: these sub-directories and inverted-files can not rollback when user drop an inverted-file index.

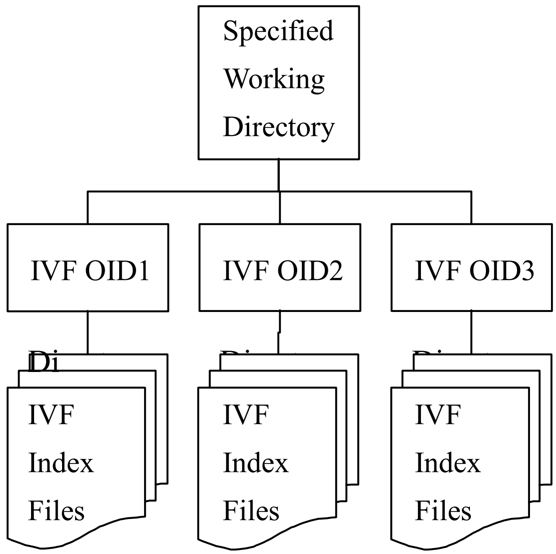

For example, \DBMaker4.1 is a specified working directory, and we create an IVF under this working directory with index name IVF1, time version 1024476670, then a sub-directory IVF1.1024476670 is created. All inverted-files are resided in the sub-directory. There are three kinds of inverted-file with different term type, Single-Byte term, Uni-Gram term and Bi-Gram term, and each inverted-file has several partitions decided by text size.

The following is the conceptual structure of IVF index files with four partitions.

Choosing between Signature and Inverted-file

Choosing between signature and inverted-file will depend on the following factors:

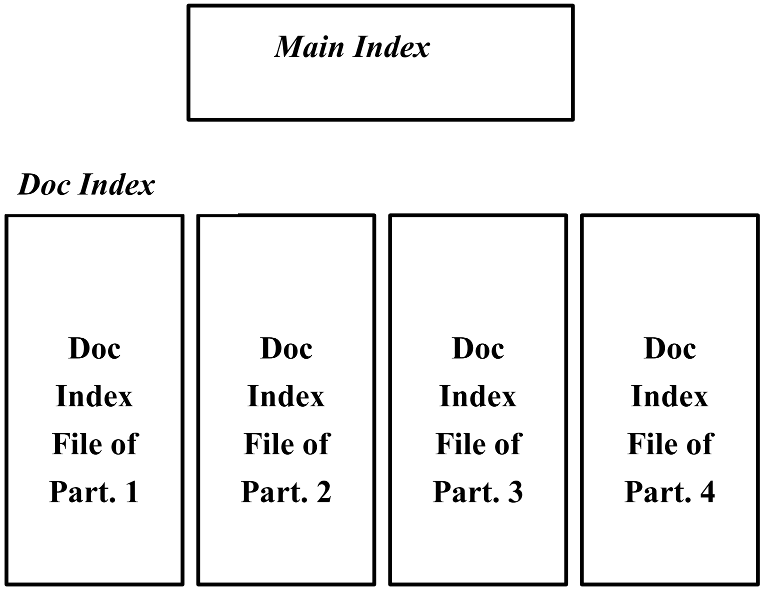

Inverted-File Structure with Four Partitions

1. Index size – the size of signature index will not exceed the ratio set by the **Scale** parameter, which is 40% of total data size by default. The avarage size of inverted-file indexes is about 1.5 times of the data size, but could grow to 2 or even 3 times, depending on the property of data.

2. Response time for queries – on a modern personal computer with sufficient memory and processing power, users can expect sub-second response time from inverted-file indexes even the data size is gigabytes. Signature indexes will take longer to respond, especially when the data size is getting large.

3. Integration with database – unlike signature indexes which are stored as BLOB objects, inverted-file indexes are stored as external files, so, for example, users cannot rollback a dropping inverted-file index operation.

Try both types of text index to find which one suits the data's characteristics best. As a rule of thumb, for data size smaller than 100 megabytes, signature indexes respond to queries reasonably fast and usually take less storage space.

Creating Text Indexes on Multiple Columns

A text index can be built using multiple columns. Use CONTAINS and the concatenation operator (||) to perform multi-column text queries. Users can query on all columns of the index or just part of them. That is, the column list in a match query must be "contained" in the column list of a text index to use the text index. Users are also allowed to use the multi-column query syntax even if no text index is created on the column list, but no text index will be used.

Searching on multiple columns is logically equivalent to merging all columns' data then searching.

➲ **Example 1**

To create an inverted-file text index **ix** on columns **author**, **subject** and **content** of the table **document**:

```
dmSQL> CREATE IVF TEXT INDEX ix ON document(author,subject,content);
dmSQL> SELECT author FORM document WHERE
    2> CONTAINS(author || subject || content, 'reagan');
```

➲ **Example 2**

To query on partial column lists:

```
dmSQL> SELECT author FORM document WHERE
    2> CONTAINS(author || content, 'reagan');
dmSQL> SELECT author FORM document WHERE CONTAINS(subject, 'reagan');
dmSQL> SELECT author FORM document WHERE subject MATCH 'reagan';
```

➲ **Example 3**

In this example, the column **subject** is included in the text index **ix** but **abstract** is not, so this query will not use any text index.

```
dmSQL> SELECT author FORM document WHERE
```

```
    2> CONTAINS(subject || abstract, 'reagan');  // no text index used
```

- **Example 4**

This example illustrates the behavior of query on multiple columns.

```
dmSQL> CREATE TABLE t1 (c1 char(20), c2 char (20), c3 serial);
dmSQL> INSERT INTO t1 VALUES('apple orange', 'banana grape');
dmSQL> INSERT INTO t1 VALUES('grape orange', null);
dmSQL> CREATE TEXT INDEX ix on t1 (c1, c2);
dmSQL> SELECT c3 FROM t1 WHERE CONTAINS (c1 || c2, 'apple');
     C3
===========
    1
 1 rows selected
dmSQL> SELECT c3 FROM t1 WHERE CONTAINS (c1 || c2, 'orange & grape');
     C3
===========
    1
    2
 2 rows selected
```

Creating Text Indexes on Media Types

DBMaker's large object columns can register the media type. For example, a LONG VARBINARY column can know its content is a Microsoft Word file, so that DBMaker can invoke proper functions to perform a full-text search on a Microsoft Word document. Table 6-1 summarizes the different media types available and associated SQL commands.

MEDIA TYPE	DATA TYPE	FILE TYPE
Microsoft Word™,	MsWordType	MsWordFileType
HTML	HtmlType	HtmlFileType
XML	XmlType	XmlFileType

Table 6-1: Media types and corresponding SQL commands

Internally, MsWordType is treated as a LONG VARBINARY object, HtmlType and XmlType are LONG VARCHAR objects, and MsWordFileType, HtmlFileType and XmlFileType are FILE objects.

Full-text Search on Media-type Columns

A user can use MATCH and CONTAINS to perform full-text search on media-type columns just as on regular text columns.

➲ **Example**

Create a table with a MS Word type column, insert some data, and search.

```
dmSQL> create table minutes(id int, doc MsWordFileType)
dmSQL> insert into minutes values(1, 'c:\meeting\20020403-1.doc');
dmSQL> select id from minutes where doc match 'Jeff';
     id
===========
    1
 1 rows selected
```

A user can also create text indexes on media-type columns.

➲ **Example**

Create a signature text index on the column **doc** of the table **minutes** and search.

```
dmSQL> create text index ix m on minutes(doc);
dmSQL> select id from minutes where doc match 'Jeff';
     id
===========
    1
 1 rows selected
```

Check Column Data's Media Type

It is possible that a media-type column contains data of different types. DBMaker will not verify the content during inserting or updating data to media-type columns; it is possible to insert a PowerPoint file to a MsWordFileType column. Altering column type to media type will change the existing data's attribute. For example, for a LONG VARBINARY column, if we alter its type to MsWordType, then the columns' data of all existing records will be considered as MsWordType thereafter, even if some of them are PowerPoint documents. If a column has data other than the designated media type, DBMaker will return a conversion error while doing match operations or creating text indexes. DBMaker provides a built-in function CHECKMEDIATYPE to check if a column's data match the column's media type. If types match, the function returns 1, otherwise returns 0.

➲ **Example**

To check where a column's data match the column's media type:

```
dmSQL> insert into minutes values(2, 'c:\meeting\20020403-1.ppt');
dmSQL> SELECT CHECKMEDIATYPE(doc) FROM minutes;

   CHECKMEDIATYPE(DOC)
========================
    1
    0
 2 rows selected
```

Dropping Text Indexes

Text indexes may be dropped using the JDBA Tool or the dmSQL DROP TEXT INDEX statement.

➲ **Example**

To drop the **ix** text index from the table **book**:

```
dmSQL> DROP TEXT INDEX ix FROM book;
```

Rebuilding Text Indexes

Unlike indexes, the text index will not simultaneously reflect table content if new records are inserted or old records are updated. Therefore, they need to be rebuilt manually. Data updated after the most recent rebuild will not be found during a text index search.

Example

```
dmSQL> CREATE TABLE song (id int, name varchar(20));
dmSQL> INSERT INTO song VALUES(1,'Endless Love');
1 rows inserted
dmSQL> CREATE TEXT INDEX ix name ON song(name);
dmSQL> INSERT INTO song VALUES(2,'Love Story');
1 rows inserted
dmSQL> SELECT * FROM song WHERE name MATCH 'love';
     id              name
=========== ====================
    1         Endless Love
 1 rows selected
```

There should be two records to match the search pattern, but only one is retrieved.

INCREMENTALLY REBUILD TEXT INDEXES

The REBUILD TEXT INDEX command rebuilds the updated data incrementally. The REBUILD TEXT INDEX command will collect all new and updated records, build new signature vectors, and append the new vectors to the tail of the text index. When only a few records have been changed, the REBUILD TEXT INDEX <*index_name*> INCREMENTAL command is the most rapid way to rebuild. The INCREMENTAL keyword is optional for incrementally rebuilding a text index.

Example

To rebuild a text index incrementally and display the results:

```
dmSQL> REBUILD TEXT INDEX ix name FOR song;
dmSQL> SELECT * FROM song WHERE name MATCH 'love';
     id              name
=========== ====================
          1 Endless Love
          2 Love Story
```

```
 2 rows selected
```

FULLY REBUILD TEXT INDEXES

If a large number of documents are deleted or updated, use the REBUILD TEXT INDEX <*index_name*> FULL command to fully rebuild a text index with its original type (signature or inverted file) and parameters.

➲ **Example 1**

To fully rebuild the **ix_name** text index for the table **song**:

```
dmSQL> REBUILD TEXT INDEX ix_name FOR song;
```

To reset the creating text index parameters of a text index or use a different type of text index, it must be dropped and re-created.

➲ **Example 2**

To rebuild the **ix_name** signature text index for the table **song** as an inverted-file index:

```
dmSQL> DROP TEXT INDEX ix_name FROM song;
dmSQL> CREATE IVF TEXT INDEX ix_name ON song(name);
```

➲ **Example 3**

To fully rebuild the **ix_name** signature text index for the table **song** with a new **total text size.**

```
dmSQL> DROP TEXT INDEX ix_name FROM song;
dmSQL> CREATE TEXT INDEX ix_name FROM song(name) TOTAL TEXT SIZE 60 MB;
```

Boolean Text Search

Not only can the MATCH operator search a simple text pattern, but also complex Boolean operations.

A user can specify the following Boolean characters in a search pattern:

'&' – AND

'|' – OR

'-' – **EXCLUDE**

'(' – **Left bracket**

')' – **Right bracket**

The precedence of Boolean characters is: bracket > EXCLUDE = AND > OR. When a MATCH pattern contains Boolean characters, all the other characters between Boolean characters are processed as simple search patterns. For example, if the MATCH pattern is "coffee | tea | apple juice", then the search pattern includes "coffee", "tea" and "apple juice".

- **Example 1**

To search for documents that contain the words 'love' and 'friend':

```
dmSQL> SELECT * FROM song WHERE name MATCH 'love & friend';
```

- **Example 2**

The following searches the documents that contain 'love' or 'friend' and does not include 'endless love'.

```
dmSQL> SELECT * FROM song WHERE name MATCH '(love | friend) - endless
love';
```

- **Example 3**

SQL syntax Boolean operators, such as AND and OR, can be used to get the same results as the MATCH pattern's Boolean operators. However, it has lower performance since only the last part of the search pattern uses the text index. For example, the following SQL command will only apply the text index scan when searching for the string 'friend', and will use a standard non-indexed search for the string 'love':

```
dmSQL> SELECT * FROM song WHERE name MATCH 'love' AND name MATCH
'friend';
```

Fuzzy Search

Sometimes users would like to search imprecise patterns. If only exact phrases are allowed, the query 'William Clinton' will not find 'William Jeffery Clinton' and vice versa. A Boolean expression like 'William & Clinton' may return many irrelevant results. DBMaker provides a fuzzy search feature that allows users to perform imprecise queries without receiving too many irrelevant results.

A phrase led by a '?' (Question mark) will be evaluated as a fuzzy expression, e.g. '?intel pentium'. Words used for a search in a fuzzy expression can be separated by up to four words in the target text. For example, `'?intel pentium'` will find `'Intel will release its 1GHz Pentium III processor'`, and `'?amd k7 athlon''AMD has renamed its K7 processor as Athlon'`.

A number of words in the query may be missing from the result set. For example `'?William Jeffery Clinton'` can find `'William Clinton'` and `'William J Clinton'`, but the first word of the query must appear; the query `'?William Clinton'` will not find `'Bill Clinton'`.

Fuzzy expressions can be combined with other text Boolean operations.

➲ **Example**

```
DmSQL> SELECT content FROM doc WHERE content MATCH '?intel pentium
& ?amd k6'

DmSQL> SELECT title FROM doc WHERE title MATCH 'al gore | ?george bush'
```

The phrase in a fuzzy expression cannot contain any other operators. Thus the expression `'?intel pentium & amd k6'` will be evaluated as `'(?intel pentium) & (amd k6)'`, and `'?(intel & pentium)'` will cause an error.

Near logic full-text search

A fuzzy match allows users to perform inexact queries without receiving many irrelevant returns. A near match search is similar to a fuzzy search, but more exact. It ensures that all words in the query string appear in the text. A phrase led by a '~'(tilde mark) will be evaluated as a near expression. For example, `'?amd sales 1ghz athlon'` will find `'AMD announced quarterly sales of its 1ghz Athlon chip'`, but not, `'AMD announced quarterly sales of its Athlon chip'`.

Near match search expressions can be combined with other text Boolean operations.

➲ **Example**

```
DmSQL> SELECT content FROM doc WHERE content MATCH '~intel pentium &
~amd k6'
DmSQL> SELECT title FROM doc WHERE title MATCH 'al gore | ~george bush'
```

Fuzzy/Near Logic Matching Rules

The following four rules apply to the matching of a query string to the result string.

1. The first word of the query must appear, e.g. the query `'?William Clinton'` will not find `'Bill Clinton'`.

2. Words can be separated by a preset number of words ("proximity"), e.g. `'?intel pentium'` will find `'Intel will release its 1GHz Pentium III processor'`, and `'?amd k7 athlon''AMD has renamed its K7 processor as Athlon'`. Currently the number of additional words in the matched result set between words in the query string can be no greater than 4.

3. A number of words in the query may be missing from the result set, e.g. `'?William Jeffery Clinton'` can find `'William Clinton'` and `'William J Clinton'`. The maximum allowed number of missing words is determined by the formula:

`max_miss = num_words - round(num_words * threshold).`

The current threshold is 0.75.

4. All words in the query must appear in the original order, e.g. `'?amd 1ghz k7 athlon'` will find `'AMD will announce 1GHz Athlon'`, but not `'AMD Athlon, formerly known as K7'`.

A phrase led by a '~'(tilde mark) will be evaluated as a near expression, e.g. `'~intel pentium'`. Near search is a special case of fuzzy search that meets the rules 1, 2 and 4 above but does not allow for missed words.

6.7 Managing Memory Tables

Memory tables, for almost all intents and purposes, function in the same manner as a regular table in DBMaker. The differences lie in the fact that memory tables are temporary tables, their life cycle being connection based. This means that a memory table exists only as long as long as there is a connection to the database. When a user severs their connection to the database, when logging out of the system to clock out for the day for example, the memory table and it's contents will be lost to the user. Memory tables are only visible during their connection to the database. Once connection to the database is lost, they are no longer visible. Unlike a regular table, memory tables are only stored in the memory of the machine that created them. They can not be used by a group and they can only have data selected or inserted, they do not support updating or delete data functions. Memory tables do support the transaction controls: commit, rollback, define save point, rollback to save point, and internal save point.

To create a memory table use the dmSQL syntax CREATE TABLE. Details on syntax and usage of the SQL command CREATE TABLE can be found in the "*SQL Command and Function Reference*".

➲ **Example**

To create memory table t1

```
create memory table t1 (c1 int, c2 char(10), c3 date);
```

Hash Index Management

Memory tables are stored in the memory of the machine that created it, for this reason memory tables do not support the B-tree structure of other table types. To help users when using memory tables users are able to create hash indexes on memory tables. Hash indexes can only be created on memory tables. The benefit of a hash index is that users have very quick access to data stored in the hash index. Hash indexes also improve equal expression and equal join performance. To create a hash index on a table users can use the CREATE HASH INDEX *index_name* ON *table_name* (*column_name*, …) [*bucket n*]; where index name is the name of the hash index being created, table name is the name of the memory table, column name is the name of the column in the memory table being effected (This value cannot specify asc/desc columns.) and bucket n sets the array size for the hash table being created. For example, with the memory table created, a hash index **inx1**, can be made on memory table **t1**, using columns **c1** and **c2** with an array size of 31.

➲ **Example**

To create hash index indx1 on memory table t1 from the previous example:

```
create hash index idx1 on t1 (c1, c2) bucket 31;
```

6.8 Managing Data Integrity

Applying constraints, or rules, to ensure the data meets certain criteria, can ensure the integrity of data. For example, verifying that an input value for a particular data item is within the correct range of values, e.g. a new employee's age must be between 16 and 90, is an example of data integrity.

In general, the different types of data integrity applicable to tables include those described in the following subsections.

Not Null

By default, all columns in a table allow NULL values. NOT NULL indicates that NULL values are not permissible in a column defined with the NOT NULL keyword.

Unique Indexes

Unique indexes, mentioned in section 6.5, "Managing Indexes", can be used to ensure no two rows of a table have duplicate values, except NULL values, in a specified column or set of columns.

Unique Constraints

A UNIQUE constraint may be set on a column, a set of columns, or an entire table. The UNIQUE constraint ensures that every row in a column has a different value. No row may have the same value in the column or columns that a UNIQUE constraint is placed on.

➲ **Example**

To create a table with the UNIQUE constraint on column Name:

```
dmSQL> CREATE TABLE student (Name CHAR(50) CONSTRAINT u UNIQUE,
                             mathematics SMALLINT);
```

Check Constraints

A CHECK constraint on a column or set of columns requires that a specified condition be true for every row of the table. If an INSERT or UPDATE statement is issued and the condition of the CHECK constraint is evaluated as false, the statement will fail.

In general, a CHECK constraint can be defined on a column (column constraint) or set of columns (table constraint).

COLUMN CONSTRAINTS

A *column constraint* is defined on a specific column and does not affect the other columns of the same table. When inserting a new row or updating an existing row, each column constraint is evaluated.

TABLE CONSTRAINTS

A *table constraint* is defined on a set of columns. When inserting a new row or updating an existing row, the table constraint is evaluated after, all column constraints are evaluated as true. Only after the table constraint is also evaluated as true will the statement be processed.

➲ **Example 1**

To create a table with column and table constraints:

```
dmSQL> CREATE TABLE student (mathematics SMALLINT
                                  CHECK VALUE >= 0 AND VALUE <= 100,
                              chemistry SMALLINT
                                  CHECK VALUE >= 0 AND VALUE <= 100,
                                  CHECK mathematics + chemistry <= 200);
```

➲ **Example 2**

To create a table with column and table constraints using standard SQL99 syntax:

```
dmSQL> CREATE TABLE student (mathematics SMALLINT
                        CONSTRAINT CHECK >= 0 AND <= 100,
                              chemistry SMALLINT
                        CONSTRAINT CHECK >= 0 AND <= 100,
                        CONSTRAINT CHECK mathematics + chemistry <=
200);
```

The keyword VALUE is used to represent the value of the column in column constraints, but the columns names are used to represent the values of the columns in a table constraint.

Primary Keys

A table can have one *primary key*, which includes a column or a group of columns with unique values to identify each row. A primary key is similar to a unique index except that its columns cannot contain NULL values. When a user creates a primary key, DBMaker will create a unique index called **PrimaryKey** on the table. After creating a table, primary keys may modified or added as long as all columns to be in the primary key contain unique, non-null values. A primary key may be added to a table or modified by using the ALTER TABLE statement. Furthermore, a primary key added to a table or modified in this fashion may be stored on a different tablespace from the table.

CREATING PRIMARY KEYS

➲ **Example 1**

To create a table with its primary key on the **ID** column:

```
dmSQL> CREATE TABLE employee (
          ID      INTEGER  PRIMARY KEY,
          name    CHAR(30),
          nation CHAR(20)
        );
```

➲ **Example 2**

To create a table with a compound primary key on the **ID** and **name** columns on tablespace **ts1**:

```
dmSQL> CREATE TABLE employee (
          ID      INTEGER,
          name    CHAR(30),
          nation CHAR(20),
          PRIMARY KEY (ID, name)
        );
```

➲ **Example 3**

To add a primary key to the **employee** table:

```
dmSQL> ALTER TABLE employee PRIMARY KEY (ID , name);
```

➲ **Example 4**

To add a primary key PK1 to the **employee** table in tablespace **TS1** using SQL99 standard syntax:

```
dmSQL> ALTER TABLE employee ADD CONSTRAINT PK1 PRIMARY KEY (ID , name)
IN TS1;
```

➲ **Example 5**

To create a table with its primary key on the **ID** column using SQL99 standard syntax:

```
dmSQL> CREATE TABLE employee (
```

```
        ID      INTEGER  CONSTRAINT pk1 PRIMARY KEY,
        name    CHAR(30),
        nation  CHAR(20)
      );
```

Dropping Primary Keys

A user can drop a primary key when it is no longer necessary. Before dropping the primary key, all foreign keys that refer to that primary key should be dropped.

➲ **Example**

To drop the primary key for the **employee** table:

```
dmSQL> ALTER TABLE employee DROP PRIMARY KEY;
```

Foreign Keys (Referential Integrity)

A column in a table containing the same values as the primary key from another table is known as a *foreign key*. A foreign key denotes the relationship between the two tables. A user can create a foreign key on a column or a group of columns in a table, and use it to reference a column or group of columns from another table. The referenced columns should be a primary key or a unique index, and cannot contain NULL values.

Referenced columns must already contain the key values being inserted into a new row for the foreign key table. If they are not present, the user will not be allowed to insert the row. In addition, all key values in the foreign key table must be deleted before deleting the key values in the referenced table.

A user can create or drop a primary key or foreign key whenever it is necessary. DBMaker will check the uniqueness of a primary key when it is created. DBMaker will also check whether all the key values already exist in the referenced table when a foreign key is created.

Creating Foreign Keys

A foreign key is used to refer to another table by specifying the referencing and referenced columns. Both the referencing and referenced columns should be mapped to each other; their schema should be the same. The mapping columns should be the same type and length. The referenced columns (specified by the primary key or unique index) should be NOT NULL, but the referencing columns (specified by the foreign key) can be NOT NULL or NULL. If the referenced column(s) are not specified, the primary key on the referenced table is regarded as the referenced column(s). Foreign keys may be created using the JDBA Tool Create foreign key wizard or the dmSQL FOREIGN KEY option.

➲ **Example 1**

To create a foreign key **f1** for the **salary** table, referencing the **employee** table with a compound primary key for its **ID** and **name** columns:

```
dmSQL> ALTER TABLE salary FOREIGN KEY f1(ID, name) REFERENCES employee;
```

➲ **Example 2**

Alternatively, specify the foreign key while creating the **salary** table:

```
dmSQL> CREATE TABLE salary (
         ID    INT,
         name  CHAR(30),
         money INT,
     FOREIGN KEY f1 (ID, name) REFERENCES employee
   );
```

➲ **Example 3**

Using SQL99 standard syntax, specify foreign key **fk1** while creating table **t1**:

```
dmSQL> CREATE TABLE t1 (
       c1 int,
       c2 int CONSTRAINT fk1 REFERENCES t2 (c1) ON DELETE SET NULL);
```

If a primary key exists for the **employee** table, a foreign key can be made for another table to refer to it without specifying the referenced columns.

Dropping Foreign Keys

If the relationship defined by a foreign key is not necessary, drop it using the JDBA Tool or the dmSQL DROP FOREIGN KEY command.

➲ **Example 1**

To drop a foreign key from the **salary** table:

```
dmSQL> ALTER TABLE salary DROP FOREIGN KEY f1;
```

6.9 Managing Serial Numbers

DBMaker provides a feature to automatically generate serial numbers. This feature is especially useful in multi-user environments for generating and returning unique sequential numbers without the overhead of disk I/O or transaction locking.

Serial numbers are signed 32-bit integers in DBMaker. A table can only have one column containing the SERIAL data type for generating serial numbers.

A user can specify the starting number for the SERIAL type column in any table when creating a table. If the starting number for a SERIAL type column is not specified, it is set to 1.

To trigger DBMaker to generate a serial number, insert a new row and supply a NULL value for the serial column. If a user inserts a new row and supplies an integer value instead of a NULL value, DBMaker will not generate a serial number. If the supplied integer value is greater than the last serial number generated, DBMaker will reset the sequence of generated serial numbers to start with the supplied integer value. SERIAL type columns cannot be defined with default values or constraints.

Creating Serial Columns

A SERIAL type column may be created using the JDBA Tool or dmSQL. A serial column must be defined with the SERIAL type keyword and an optional starting number.

➲ **Example**

To create a SERIAL type data column **ID** for the **employee** table and specify its starting number as 1001:

```
dmSQL> CREATE TABLE employee (nation CHAR(20) DEFAULT 'R.O.C',
                                        ID SERIAL(1001),
```

```
                                name CHAR(30) NOT NULL,
                            joinDate DATE DEFAULT CURDATE(),
                              height FLOAT,
                              degree VARCHAR(200)) IN ts1;
```

GENERATING SERIAL NUMBERS

DBMaker automatically generates serial numbers when a NULL is inserted into the SERIAL type column.

➲ **Example**

To insert a new row into the **employee** table and generate a serial number for the column **ID**:

```
dmSQL> INSERT INTO employee VALUES
              ('U.S.A', NULL, 'Jeff', , 6.6, 'Director', NULL);
```

RETRIEVING SERIAL NUMBERS

DBMaker keeps the last generated serial number in the LAST_SERIAL column of the system table SYSCONINFO for each connection. After inserting a record containing a serial number, the serial number can be retrieved from LAST_SERIAL.

➲ **Example**

To get the serial number that has just been generated for the inserted record:

```
dmSQL> select LAST SERIAL from SYSCONINFO;
LAST SERIAL
===========
        200
 1 rows selected
```

RESETTING SERIAL NUMBERS

A user can reset the counter for a serial column. This allows a new sequence to be started in a serial column without having to modify the table.

➲ **Example**

To alter the serial counter value for the **employee** table from its current value to 3000:

```
dmSQL> ALTER TABLE employee SET SERIAL 3000;
```

6.10 Managing Domains

A *domain* is a type of integrity constraint used when defining a column. Domains specify the data type for the column, and may specify a default value or a value constraint. When a column is defined using a domain, it inherits the properties of the domain, (data type, default value, and value constraint), without requiring the user to specify them.

Specifying the default value and value constraint using domains achieves the same result as specifying them in a standard column definition. If a user specifies a default value for a column, it will override the default value specified in a domain.

Any value constraints specified in the column definition will be used in addition to the value constraints specified in the domain. If a user defines a column using a domain and specifies additional value constraints, the additional value constraints must not conflict with those defined in the domain.

DBMaker does not check for conflicting value constraints, so it may be possible to define value constraints that would not allow the user to enter any values at all. All data types supported by DBMaker except the SERIAL type can be used in domains.

Creating Domains

A domain is defined by a domain name, an optional default value, and an optional constraint. For example, a user might want to ensure that all columns dealing with some form of titles, (movie, CD, or videotape), have a data type of VARCHAR, are no more than 35 characters in length, and do not permit insertion of NULL values. Domains may be created using the JDBA Tool or the dmSQL CREATE DOMAIN statement.

➲ **Example 1**

The keyword VALUE is used to represent the value of the column defined on the domain. To create a specific domain that is used in the subsequent CREATE TABLE statements:

```
dmSQL> CREATE DOMAIN title_type VARCHAR(35) CHECK VALUE IS NOT NULL;
```

➲ **Example 2**

To define columns as in the CREATE TABLE statement:

```
dmSQL> CREATE TABLE movie tit                                    les (title
title_type, ..., ...)
```

DROPPING DOMAINS

A domain can be dropped only when there are no columns referenced on it. Domains may be dropped using the JDBA Tool or the dmSQL DROP DOMAIN statement.

➲ **Example**

To use the DROP DOMAIN statement:

```
dmSQL> DROP DOMAIN title_type;
```

6.11 Unloading / Loading Objects

Sometimes the user may need to save database data to an external text file. DBMaker provides the UNLOAD and LOAD commands just for this purpose. Objects that are unloaded from the database are not removed from the database; they are simply saved as one or more external text files. When an object is loaded onto a database, the schema of that object is also recreated.

Unloading Objects

Unload is a tool provided by dmSQL used to transfer the contents of a database to an external text file. After the unload procedure succeeds, dmSQL will produce two text files. One stores the script, with extension name **.s0**, to establish the database object and the other stores the BLOB data, with the extension name **.bn**.

There are eight options for the unload command: unload database, unload table, unload schema, unload data, unload project, unload module, unload procedure, and unload procedure definition. Unloading an object requires that the user have SELECT privilege on the object in question. For instance, if a user has SELECT privilege on a table, then only that user can unload the content of this table. Only a user with DBA or SYSADM authority may unload the database.

UNLOAD [DB | DATABASE]

A user with DBA or SYSADM authority may unload the content of a database to an external text file. This file includes information about security, tablespaces, definitions, indices, synonyms, data, etc. For each database, dmSQL will generate at least two external files, one script, and one BLOB data.

➲ **Example**

```
dmSQL> unload db to empdb;
```

The name of the external text file is empdb. By default, dmSQL will create these files in the current working directory. In the above statement, there are at least two text files created, **empdb.s0** and **empdb.b0**. If the unloaded BLOB file **empdb.b0** exceeds the maximum size allowed by the operating system, dmSQL will sequentially generate files **empdb.b1**, **empdb.b2**, …, **empdb.bn**, up to a maximum number of 99. dmSQL will always generate one script file **emodb.s0**, with a maximum size limited by the operating system.

UNLOAD TABLE

The UNLOAD TABLE command unloads tables to an external file and will record the definition, synonyms, indices, primary key, foreign keys, and data of the table. Use the wild cards "_" and "%", which correspond to "?" and "*" in DOS in the owner and table name. The wild card "_" represents a character, and "%" represents a set of characters.

UNLOAD SCHEMA

The usage of this option is very similar to unload table. It can only unload the definition of a table; it cannot unload the data in a table. It uses the same wild cards as the UNLOAD TABLE option.

UNLOAD DATA

This option will unload all data from a table. It will not unload the definition of the table. UNLOAD DATA uses the same wildcards as the previous two options. Only users with the SELECT privilege on the unloaded table may execute the UNLOAD DATA command.

DBMaker 3.6 and later versions support an additional syntax for unloading data:

```
dm SQL>unload data from (select statement) to file_name.
```

If the select statement is a join, the projection columns must be from the same table, the following statement is executable. DDL commands, delete, insert, or updates are not permitted.

➲ **Example 1**

Valid syntax:

```
dmSQL> unload data from (select t1.c1, t1.c2 from t1, t2 where t1.c1=
t2.c1) to f1;
```

➲ **Example 2**

Illegal syntax:

```
dmSQL> unload data from (select t1.c1, t2.c1 from t1, t2 where t1.c1 =
t2.c1) to f1;
```

No aggregate or built-in functions are permitted in the projection columns.

➲ **Example 3**

Illegal syntax:

```
dmSQL> unload data from (select avg(c1) from t1) to f1;
dmSQL> unload data from (select now() from t1) to f1;
```

Views and synonyms are permitted.

➲ **Example 4**

Valid syntax:

```
dmSQL> unload data from (select * from s1 where c1 > 10) to f1;
dmSQL> unload data from (select * from v1 where c1 < 10) to f1;
```

UNLOAD PROJECT

This option allows a user to unload an ESQL/C project to an external text file.

UNLOAD MODULE

This option allows a user to unload a module to an external file.

UNLOAD [PROC | PROCEDURE]

This option allows a user to unload the stored procedures to an external file.

UNLOAD [PROC DEFINITION | PROCEDURE DEFINITION]

This option allows a user to unload the definition of the stored procedure to an external text file.

➲ **Example 1**

The following will unload the table **e tab** for the current user; if there are any blanks in the table name add double quotes:

```
dmSQL> unload table from "e tab" to empfile;
```

➲ **Example 2**

The following will unload all tables with the names starting with **emp** for the SYSADM owner, for example, **emptab**, **empname**, … etc:

```
dmSQL> unload table from SYSADM.emp% to empfile;
```

➲ **Example 3**

The following will unload the schema of all tables with the name **ktab**:

```
dmSQL> unload schema from %.ktab to kfile;
```

Unload the table with names containing wild cards. Use the escape character “\”, or double quotes on the name.

➲ **Example 4**

The following commands will unload data from a table named **abc%**:

```
dmSQL> unload data from abc\% to abcfile;
dmSQL> unload data from “abc%” to abcfile;
```

Loading Objects

The LOAD command is a tool provided by dmSQL, and is used to transfer a database object that has already been unloaded to a text file, into the database. There are seven options: load database, load table, load schema, load data, load project, load module, and load procedure. A file must be unloaded and loaded with the same option. For example, load a database from a text file that was unloaded with the database option. When loading a text file, set the number of commands <*n*> to automatically commit the transaction. The default number is 1000. The size of <*n*> will affect whether the transaction succeeds or not and the loading speed. The journal will fill easily with a large <*n*> value and could cause the transaction to fail. A small <*n*> value will increase the number of transactions committed and slow down the loading speed. If there are errors occurring during the loading procedure, an error messages will be recorded in a log file, which the system will use to undo executed commands. The log file is stored in the same directory as the external text file being loaded and does not stop the loading procedure.

LOAD [DB | DATABASE]

Use the LOAD [DB | DATABASE] command to transfer the contents of a database to a new database. First, unload the database to transfer to an external text file, and then use the LOAD DB command to load the contents of the database from the text file. Before loading a database, create a new one. The name of the new database can be different from the old one. Only a DBA or a SYSADM may execute this command.

Example

The following set option for LOAD DB has been added to versions above DBMaker 3.6:

```
Set LOAD DB [safe | fast]
```

The database will run in normal mode if LOAD DB is set to SAFE. The load utility will rollback to the last committed command if any error occur during loading, the error messages will return to screen, and write to the log file of the load utility. When using the set LOAD DB in fast mode, the rule for loading the utility in DBMaker versions earlier than 3.6, will make the whole load procedure work under the no journal mode. Setting LOAD DB in fast mode will speed up the load utility, but it will make the database shut down in no journal mode if any error occurs. For example, suppose that the load file has tablespace creation but it is not specified in the dmconfig.ini file. If LOAD DB is set to use the safe option, the following error message, "ERROR(8002): [DBMaker] keyword entry is required for configuration file", will be reported and then the load command will rollback. If LOAD DB is set to use the fast option, then the following error message occurs, "ERROR(30017), [DBMaker] errors occurred in no-journal mode, shut down database".

NOTE *The default option is SET LOADDB SAFE".*

LOAD TABLE

This option permits loading the contents of a table, including schema and data, from a text file. When loading a table from a text file, make sure that the table name is unique.

LOAD SCHEMA

This option allows users to load the schema, not including the data, from a table contained in a text file. When loading a table schema from a text file, ensure that the table name is unique.

LOAD DATA

A corresponding table must exist when loading data from an external text file. In versions earlier than 3.6 when the errors occur during the LOAD DATA procedure, it will rollback to the last committed command.

➲ **Example**

DBMaker 3.6 and later versions support the following options:

```
Set load data skip [error] | stop [on error]
```

If LOAD DATA SKIP ERROR is set then the following error messages will be skipped during the loading of data:

ERROR(401) unique key violation.

ERROR(410) referential constraint violation: value does not exist in parent key.

ERROR(6521) table or view does not exist.

ERROR(6002) syntax error.

ERROR(6015) incomplete SQL statement input.

The errors will be skipped and the load utility will resume execution of subsequent commands. The above errors are the most common errors to occur during loading of data. When LOAD DATA STOP or STOP ON ERROR is set, the whole LOAD command will be rolled back if an error occurs. The default value for this option is LOAD DATA SKIP ERROR. All the error messages that occur during the loading of data will be written into the log file.

LOAD MODULE

This option allows a user to load a module from an external text file.

LOAD PROJECT

This option allows a user to load an ESQL/C project from an external text file.

LOAD [PROC | PROCEDURE]

This option allows a user to load a stored procedure from an external text file.

Example 1

The following command loads the database from a file named **empdb**, and commits it automatically every 100 commands during loading. The system will generate a log file named **empdb log** in the same directory:

```
dmSQL> load db from empdb 100;
```

Example 2

The following command will load a table from a file named **empfile**, and it will commit automatically every 50 commands during loading:

```
dmSQL> load table from empfile 50;
```

Example 3

The following command will permit the loading of data from an external data file named **datafile** and will commit automatically every 1000 commands using the default setting:

```
dmSQL> load data from datafile;
```

6.12 Browsing System Catalogs

DBMaker keeps detailed information on all schema objects in the *system catalog* tables. For more information on system catalog tables, see "Appendix B".

SCHEMA OBJECT INFORMATION	SYSTEM CATALOG TABLE NAME
Tables	SYSTABLE
Columns	SYSCOLUMN
Views	SYSVIEWDATA
Synonyms	SYSSYNONYM
Indexes	SYSINDEX

Domains	SYSDOMAIN
Serial numbers	SYSCONINFO
Table constraints	SYSTABLE
Column constraints	SYSCOLUMN
Domain constraints	SYSDOMAIN

Table 6-2: Schema information in the System Catalog tables

6.13 Calculating the Space Required

As stated in previous sections, only tables and indexes occupy physical disk space. To manage disk space, estimate the size of each object and decide which tablespace each object will belong to before creating them. In the estimation phase, the database administrator must have a clear picture of how to construct tablespaces using tables and how much hardware will be required to support the database in the future.

Generally, tables that are split between several tablespaces will get higher performance than tables in a single large tablespace. Conversely, many small tablespaces are harder to manage.

How to Estimate the Size of a Table

The following formulas show how to estimate the size of a table and the size of an index:

table size = row size × number of rows × 1.05

index size = key size × number of rows × 1.20

These two formulas are used to estimate the size needed for a tablespace by adding the size of all tables and indexes in it. In the above formulas, 1.05 and 1.20 are estimates of the resource overhead used to calculate the system resources required. The row size and key size contain the internal record header size. The following subsections show how to calculate the size of a row and a key.

Row Size

The storage size of a row, excluding BLOB data, cannot exceed 3996 bytes and consists of the space required for data storage and an internal record header.

The size of an internal record header is equal to:

internal record header size = (no. of columns + 1) × 2

Each data type has space requirements:

TYPE	COLUMN LENGTH
BINARY(n)	n
CHAR(n)	n
SMALLINT	2
INTEGER	4
FLOAT	4
SERIAL	4
DOUBLE	8
DECIMAL(p,s)	[(p+1)/2]+2
TIME	4
DATE	4
TIMESTAMP	11
OID	8
VARCHAR(n)	0.n
FILE	20
LONG VARBINARY	20+X
LONG VARCHAR	20+X

Table 6-4: Data Types and sizes

NOTE *The size of a VARCHAR type column depends on the size of the actual data, but cannot exceed the maximum size defined. A BLOB type column (LONG VARCHAR or LONG VARBINARY) will occupy at least 20 bytes in the data file* *and the actual data will be stored in the BLOB file or a data file. For more detailed information, refer to Chapter 7, "Large Object Management". If the value in a column is NULL, it does not occupy any space.*

➲ **Example**

To create a table with five columns defined:

```
dmSQL> CREATE TABLE employee (ID INTEGER NOT NULL,
                            name CHAR(30) NOT NULL,
                          height FLOAT,
                          degree VARCHAR(200),
                         picture LONG VARCHAR);
```

After issuing this command, insert rows into the table and calculate the size of the record:

(3001, "Jeff Yang", 175.5, "Stanford PhD.", [pic1]) where **pic1** is an image.

DATA ITEM	TYPE	SIZE
ID	integer	4 bytes
name	char	30 bytes
height	float	4 bytes
degree	varchar	13 bytes
picture	long varchar	20 bytes
row header	—	12 bytes
	Total	**83 bytes**

= (4 + 30 + 4 + 13 + 20) + (5 + 1) × 2

= 83 bytes

(3002, "George Wang", 180.0, "NCTU Ms.", NULL)

DATA ITEM	TYPE	SIZE

DATA ITEM	TYPE	SIZE
ID	integer	4 bytes
name	char	30 bytes
height	float	4 bytes
degree	varchar	8 bytes
picture	long varchar	0 bytes
row header	—	12 bytes
	Total	**58 bytes**

= (4 + 30 + 4 + 8 + 0) + (5 + 1) × 2

= 58 bytes

DBMaker will verify that the row size does not exceed 3996 bytes when inserting or updating rows. When creating a table, DBMaker also verifies that the smallest possible row size does not exceed 3996 bytes.

The smallest row size in the above example can be calculated as follows:

minimum row size = (4 + 30 + 0 + 0 + 0) + (5 + 1) × 2

= 46 bytes

The minimum row size does not exceed 3996 bytes, so DBMaker will allow this table to be created.

KEY SIZE

Key storage size is made up of the space required for data storage in the index columns and an internal record header. It also requires an extra eight bytes for an internal row identifier. The internal row identifier also requires one byte in the record header.

The size of the internal record header is equal to:

Internal record header size = (no. of columns + 1 + 1) × 2

For example, if an index is created on a single column with the SMALLINT type, the size of each key will be:

key size = 2 + 8 + (1 + 1 + 2) × 8

= 26 bytes

In this case, two bytes are used by the data in the key column, eight bytes are used for the internal row ID for each key, and six bytes are used for the record header.

Estimating the Size of Tablespaces and Tables

The following example demonstrates how to estimate the size of a tablespace and its tables. Assume there is a tablespace that contains three tables **A**, **B**, and **C**, and one index **D** created for table **A**. Columns in table **A** are defined as INTEGER, and CHAR(10). Columns in table **B** are defined as SMALLINT, CHAR(10), FLOAT, and VARCHAR(200). Columns in table **C** are defined as SMALLINT, INTEGER, and LONG VARCHAR. Index **D** is created on the first column in table **A**. Table **A**, table **B** and table **C** consist of 1500, 3000, and 250 rows respectively.

The row and key sizes for this database can be calculated as shown below. Suppose the average length of the VARCHAR column in table **B** is 80 bytes, and the size of each BLOB column in the data file in table **C** is 20 bytes:

In table **A**: row size = (4 + 4 + 10) + 8
= 26 bytes

In table **B**: row size = (2 + 10 + 4 + 80) + 10
= 106 bytes

In table **C**: row size = (2 + 4 + 20) + 8
= 34 bytes

If the average size of each BLOB item in table **C** is 9000 bytes, then specify the BLOB frame size to be 11KB. See Chapter 7, "Large Object Management" for more information about BLOB data.

Index **D**: keysize = 4 + 4 + 8
=16 bytes

The table sizes for this database can be calculated as shown below. Note that the size of table **A** also includes the size of index **D**.

Table **A**:	table size = (26 × 1500 × 1.05) + (16 × 1500 × 1.2)
	= 40950 + 24000 bytes
	= 10 + 6 pages
	= 16 pages
Table **B**:	table size = 106 × 3000 × 1.05
	= 333900 bytes
	= 82 pages
Table **C**:	table size = 34 × 250 × 1.05
	= 34 × 250 × 1.05 bytes
	= 3 pages

In the BLOB file, the size of table C is 250 frames (every row needs a frame).

After examining the above figures, create a tablespace with at least one data file (16+82+3=101 pages) and one BLOB file (250 frames with a frame size of 11KB) to store the above tables and index.

Estimate the size of a tablespace when creating it to avoid the trouble of adding or enlarging files later.

6.14 Checking Database Consistency

DBMaker includes several commands that a user with DBA privilege can use to check the consistency of a database. Examples of database consistency include an index that has a key but does not exist in the table, or a key that exists in a foreign table but does not exist in the parent table. DBMaker supports six commands to check different levels of consistency. These commands are time consuming when the database is large and they will take locks, administrators should only use them when necessary.

Checking Indexes

DBMaker allows a user with privilege on the index in question to check an index and its relationship to a table. It checks if the index structure (B-tree) is correct, if the data is in order, and if the index keys exactly match the data records, etc.

If an index seems to have a problem, use this command to verify that a problem exists. If DBMaker finds inconsistencies in the index, drop and rebuild the index to fix it.

➲ **Example**

To check the index consistency for the index **IDX1** in the **TBL1** table:

```
dmSQL> CHECK INDEX TBL1.IDX1;
```

Checking Tables

DBMaker allows a user to check all records, indexes, and BLOB data associated with a table and the relationship between foreign and parent tables, given that the user has privilege on those objects. If there is any inconsistency in a table, unload all records in the table, drop the table, recreate it, and then reinsert all records.

➲ **Example**

To check consistency for the **TBL1** table:

```
dmSQL> CHECK TABLE TBL1;
```

Checking Catalogs

DBMaker allows a user with DBA privilege to check the consistency of system tables. If the system catalogs have errors, the database may be seriously corrupted.

➲ **Example**

To check the consistency of the system catalogs:

```
dmSQL> CHECK CATALOG;
```

Checking Databases

DBMaker also allows a user with DBA privilege to check the whole database including the system catalogs and all tablespaces.

➲ **Example**

To check the consistency of an entire database:

```
dmSQL> CHECK DB;
```

If corruption exists and the database has been backed up, use the most recent backup to restore it. For more information, refer to Chapter 14, *Database Recovery, Backup, and Restoration*.

When the database has no backup and an index is corrupted, drop and recreate the affected indexes. If any other type of corruption has occurred, immediately backup the database including all data and journal files. Then try to shut down and restart the database, then run the CHECK commands again. After DBMaker automatically recovers from a crash, some types of corruption may be fixed. If any inconsistency still exists, contact a CASEMaker technical support representative to help fix the remaining problems with the database.

6.15 Updating Statistics for Schema Objects

Outdated statistics values for schema objects (tables, indexes, columns) may cause the DBMaker optimizer to use an inefficient plan for an SQL statement. If users have inserted large amounts of data into the database after the last time the database administrator updated the statistics values, update the values again.

➲ **Example 1**

To update the statistics values for all schema objects:

```
dmSQL> update statistics;
```

If a database is extremely large, it will take a lot of time to update statistical values for all of the schema objects. An alternative method is to update statistics on specific schema objects that have been modified since the last update, and set the sampling rate.

➲ **Example 2**

To update statistics for tables:

```
dmSQL> update statistics table1, table2, user1.table3;
```

➲ **Example 3**

To update statistics for index **idx1** on **table1**:

```
dmSQL> update statistics table1 (index idx1);
```

➲ **Example 4**

To update statistics for tablespace **ts1**:

```
dmSQL> update tablespace statistics ts1;
```

➲ **Example 5**

To update statistics for indexes idx1 and idx2 on table1:

```
dmSQL> update statistics table1 (index idx1, idx2);
```

Please refer to Chapter 17, *Performance Tuning* for more information about updating statistics and the SQL optimizer.

7 Large Object Management

A *Large Object* (LO) is any variable length data object, such as document text, images, sound, or video. DBMaker has a great deal of flexibility when dealing with large objects and provides an excellent mechanism for unstructured data.

DBMaker does not limit the number of LOs that can be in a table, and there is no aggregate size limit on LO columns. This means the capacity of each LO column can be up to 8TB. DBMaker can use extensions to the SQL language to directly access Large Objects, eliminating the need for users to learn any special syntax. All access to LO columns is transparent in SQL statements, which makes using large objects easy to learn. Furthermore, users can input or output LO data to and from a file using SQL commands or the ODBC interface.

If the same LO data is contained in many tuples, DBMaker will store only a single LO and share it between the tuples instead of duplicating a copy for each tuple. This can decrease disk utilization dramatically. However, from the user's view there is always a dedicated LO item for each tuple. For example, if one user updates a shared LO in one tuple, the other tuples that share that LO are not influenced and users still see the unchanged LO data. The LO that was updated will be stored as a new LO in the database.

An LO is always written to disk as a single unit. However, users can read all or part of an LO. The SELECT, UPDATE, INSERT and DELETE statements are permitted with LOs. LO items can only be used in Boolean expressions if users would like to test them for NULL values. DBMaker also provides the MATCH function for use with LOs to perform searches with pattern matching. The MATCH function is similar to the LIKE function except it only works on LO columns and does not permit the use of wildcard characters.

DBMaker does not permit the operation of arithmetic or string expressions on LO items, nor can the LO items be used in any of the following ways:

- With aggregate functions
- With the IN, ANY, EXIST or LIKE predicates
- With the GROUP BY clause
- With the ORDER BY clause

There are two kinds of LOs: Binary Large OBjects (BLOBs), which are stored in database files, and File Objects (FOs), which are stored as external files on a host file system.

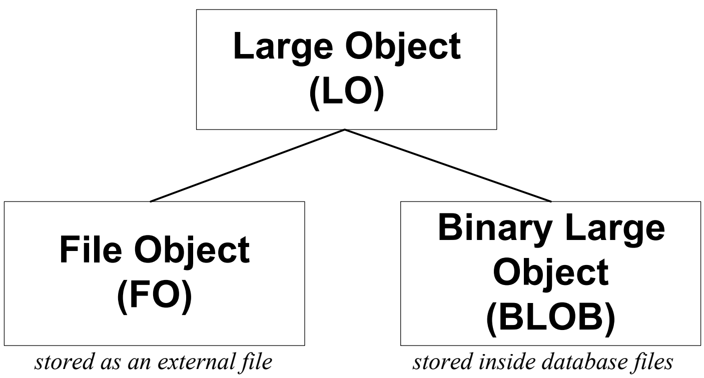

Figure 7-1: Large Objects supported by DBMaker

A BLOB, stored in database files, can only be accessed through DBMaker and insists on the data integrity provided by DBMaker, such as transaction controls, logging and recovery. A BLOB can only be shared between tuples in the same table. However, a FO can be shared between tables in a database. In addition, when the data needs to be shared by the other non-database applications, using FOs will be more flexible.

7.1 Managing BLOBs

There are two types of BLOB items, LONG VARCHAR (or CLOB) and LONG VARBINARY. Data of the LONG VARCHAR type can consist of any kind of text data such as memos, long text, HTML source files, or program source files. Data of the LONG VARBINARY type can be any kind of binary data such as images, sound, spreadsheets, program modules, etc.

A BLOB may be stored in a data file or a BLOB file, depending on its size. Although the format of a data file is fixed, the format of BLOB files in the database should be customized to obtain better performance and disk utilization.

The choice of BLOB logging is optional because it occupies a large amount of the journal space and can pull down performance. To save logging space and improve performance, the BLOB journal may be turned off. However, if BLOB logging is turned off, DBMaker will not ensure that the BLOB contents will be correct after the database has been restored from a backup. If the BLOB journal is turned on, make sure the journal file has enough space to accommodate the BLOB data.

Customizing BLOB Space

DBMaker automatically decides where to store BLOB data. If the size of a LONG VARCHAR or LONG VARBINARY column is small and the total length of a tuple does not exceed the limitation for the maximum tuple size, DBMaker will put the BLOB data in a column together with the other data in the database. This increases efficiency because the BLOB data is also fetched when DBMaker fetches a tuple.

When the total length of a data tuple exceeds the limitation of the maximum tuple size, DBMaker will store the BLOB data separately. In this situation, getting the BLOB data (called an indirect BLOB) requires two disk operations, one to fetch the data tuple, and one to fetch the BLOB data.

According to its size, an indirect BLOB may be stored in a DATA file or in a BLOB file in the same tablespace as the table. The data in an indirect BLOB column is stored in a data file when its size is equal to or less than 3950 bytes. Otherwise, it is stored in a BLOB file.

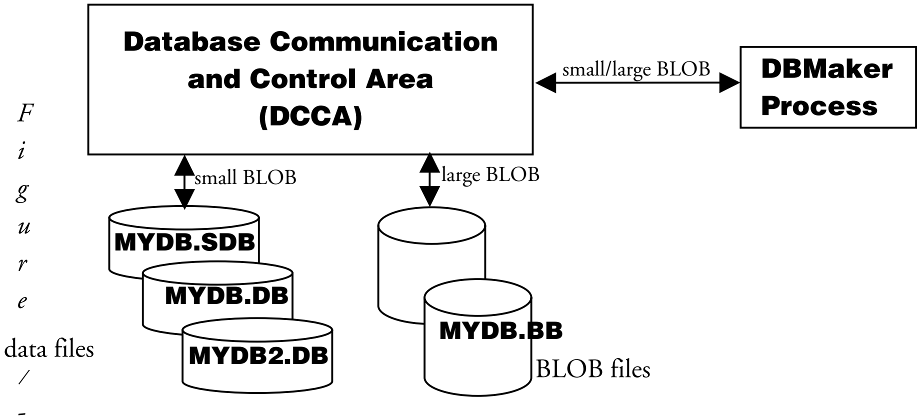

Figure 7-7-2: DBMaker accesses BLOB data through DCCA

Data files contain pages, and BLOB files contain frames. There are two major differences between pages and frames:

- The size of a page is fixed at 4KB, but a user can customize the size of a frame.
- A page can contain more than one tuple, but a frame can only contain a single BLOB.

The frame size of a BLOB file can be customized before database creation to increase performance and disk utilization. To customize the frame size, specify the value in kilobytes of the DB_BFRSZ configuration keyword in **dmconfig.ini**. The default value of DB_BFRSZ is 16. Refer to Section 4.2, *Creating a Database* for more information on configuration parameters that must be set before database creation.

➲ **Example**

To specify the BLOB frame size by adding a line to the **dmconfig.ini** file:

```
DB_BFRSZ = 10                              ; BLOB frame size = 10K bytes
```

The valid range of DB_BFRSZ is 8 to 256, except in the Microsoft Windows 3.1 environment where the frame size is fixed at 8KB.

The frame size of all BLOB files in a database is the same. Once a database is created, the BLOB frame size cannot be changed from its initial setting. DBMaker will keep this value in the database system information table. When the database is restarted, DBMaker will get the original value from the system information page and ignore the DB_BFRSZ keyword in **dmconfig.ini**.

➲ **Example**

To query the SYSINFO system table for the frame size:

```
dmSQL> SELECT INFO, VALUE FROM SYSINFO WHERE INFO = 'FRAME SIZE';

              INFO                             VALUE
================================ ================================
FRAME SIZE                       16384

1 rows selected
```

Determining the frame size is a trade-off between disk utilization and performance. If entire BLOBs are frequently retrieved, adjusting the frame size to contain the entire BLOB will result in better performance because only one disk access is required. However, there may be large variations in the size of the BLOB data. If the frame size is set large enough to contain the largest BLOB, it may waste disk space, as other frames that contain smaller BLOBs will contain unused disk space.

Alternatively, the frame size is only large enough to contain the smallest BLOBs, performance will be degraded when fetching larger BLOBs that are stored in multiple frames.

Each frame contains a header to record frame information. If the frame size is 8KB, for example, the space occupied by the BLOB will be less than 8192 bytes. About 1.8KB is reserved to store information (such as where other frames are) for each BLOB item, so the usable space in the first frame of a BLOB is much less than the size of the entire frame. Thus, if the actual size of a BLOB is 8192 bytes, it will occupy two frames: the first 6.2KB of the BLOB are stored in the first frame and the remaining bytes of the BLOB are stored in the second frame.

A group contains 252 frames. The first frame of the group is a 4-KB directory page. The remaining 251 frames are for data, and their size is determined by DB_BFRSZ.

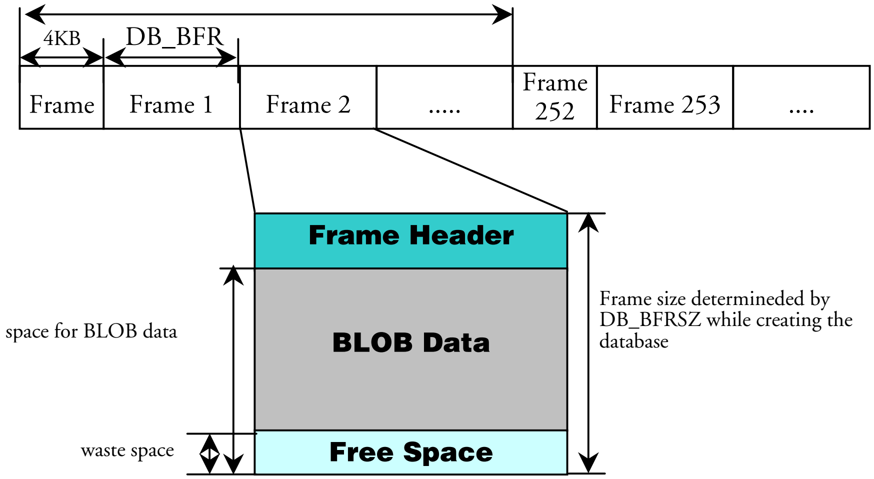

Figure 7-3: the structure of a BLOB file

Users can calculate the size of a BLOB file according to the total number of frames, directory pages and data frames.

The following formula shows how to calculate the size of a BLOB in KB:

Number of directory pages = total frames/252

BLOB frames = total frames – number of directory pages

BLOB file size = number of directory pages x 4 KB + BLOB frames x DB_BFRSZ

For example, if the BLOB frame size is 16-KB, the size of a BLOB file with five frames is:

1 x directory page + 4 x data frames

= 1 x 4KB + 4 x 16KB

= 68 KB

DB_BbFil specifies the system BLOB file name in the system tablespace, SYSTABLESPACE. Users cannot specify the size of the system BLOB file. The default file name for the system BLOB file is the database name concatenated with '.SBB' 。

DB_UsrBb specifies the default user BLOB file name in the default tablespace, DEFTABLESPACE, and its size.

For more details on adding new BLOB files to an existing user tablespace, refer to the subsection *Adding Files to Tablespaces* in Section 5.3.

Generating BLOBs

A BLOB column is the same as other columns except that its data type is LONG VARCHAR or LONG VARBINARY.

➲ **Example**

To create two BLOB columns named **note** and **photo**:

```
dmSQL> CREATE TABLE employee (id INTEGER, note LONG VARCHAR, photo LONG
VARBINARY);
```

Insert BLOB data from the **aa.txt** file and an image file **img001.jpg**:

```
dmSQL> INSERT INTO employee VALUES(1, 'aa.txt', 'img001.jpg');
```

Alternatively, insert BLOB data from the **ab.txt** file and image file **img001.gif** using host variables:

```
dmSQL> INSERT INTO employee VALUES(2,?,?);
dmSQL/Val> &ab.txt, &img001.gif(2,4);
dmSQL/Val> END;
```

The resulting LONG VARBINARY column is represented in hexadecimal format. The following results will be returned when browsing the table:

```
dmSQL> SELECT * FROM employee;
 id          note              photo
=====      ===========     =================
   1       <!-- TWO ST     474946383961b705d
   2       <script lan     ffd8ffe000104a464
```

DBMaker also supports fetching BLOB data into a user-specified file. For more information on how to insert and fetch BLOB data, refer to the "*JDBA Tool User's Guide"* and the *"ODBC Programmer's Guide*".

Updating BLOBs

A BLOB item is always written to disk as a whole. Thus, when updating a BLOB column, DBMaker will drop the original BLOB item and then insert the new data as a new BLOB item.

Example

To update contents for a BLOB column, using the UPDATE command:

```
dmSQL> UPDATE employee SET note = 'Hello !' WHERE id > 0;
dmSQL> SELECT * FROM employee;
 id          note               photo
=====      ===========      =================
   1       Hello !          31323334353637
   2       Hello !          33343536
```

From the user's viewpoint, there must be a BLOB for each tuple. However, to save disk space, DBMaker only creates a single copy of the BLOB data shared by all tuples with an ID greater than zero. DBMaker maintains an internal counter to record how many tuples reference a BLOB. When updating a BLOB column for a tuple that links to the shared BLOB, DBMaker will generate a new BLOB item and decrease the counter of the shared BLOB by a value of one. This prevents any changes made to the BLOB column from influencing other tuples. In DBMaker, this is known as *loose coupling*. This makes disk utilization more efficient, but a BLOB item can only be shared by tuples that are in the same table. If a BLOB is not linked with tuples, DBMaker will automatically drop it.

Predicate Operations on BLOB Columns

BLOB objects can only be used in CONTAIN, MATCH or Boolean expressions when testing for NULL values.

Example 1

To fetch all data from the "employee" table from the **NOT NULL** "note" column:

```
dmSQL> SELECT * FROM employee WHERE note IS NOT NULL;
```

DBMaker provides pattern matching for BLOBs. The CONTAIN and MATCH function is similar to the LIKE function except that wildcard characters are not supported. The difference between CONTAIN and MATCH is that the former is a partial word match and the latter is a full word match. For example, 'This is a character.' CONTAIN 'char' and 'This is a character.' MATCH 'character', but 'This is a character.' NOT MATCH 'char'.

➲ **Example 2**

To find all employees from the "notes" column containing 'Database Administrator':

```
dmSQL> SELECT * FROM employee WHERE note MATCH 'Database Administrator';
```

7.2 Managing File Objects

Each *file object* (FO) column references external files. Using FOs is beneficial when the data is also used by other applications, since the file can be accessed directly. Most current multimedia tools can only process multimedia data when it is stored as a complete file of the required type. Multimedia data that is stored in BLOB or data files must be fetched from DBMaker by the user and redirected to a file that can be processed by the appropriate tool. However, if BLOB data is stored as an FO a user can simply get the file name from DBMaker and pass the name to the appropriate multimedia tool.

There are two kinds of FOs: *system* and *user*. System file objects are created when a user inserts data on the client side and DBMaker passes it through and stores it in an external file specified by the configuration parameter **DB_FoDir**. System FOs are created by DBMaker and can be recognized by their default .FOB file name extension. User file objects are external files that are simply linked to a column. A user file object may be a file on any device that is accessible by DBMaker through the server's operating system.

The major difference between system and user file objects is that DBMaker generates a system FO automatically. This means a system FO will be deleted when no column references it. Therefore, with system FOs, users can leave storage management to DBMaker. Another advantage to using system FOs is that data is duplicated to the server side, so users can manage data from the server. DBMaker's backup and restoration features also protect system FOs.

A user FO will not be deleted when there are no more references to it. The major advantage to user FOs is that DBMaker can link a column to an existing file directly. It does not need to duplicate data, such as a file on a CD-ROM. This conserves disk space and makes it easier to share a file among several records. However, if a file is deleted outside of DBMaker, all columns linking to this file may become invalid. A file linked as a user FO must open its read permission.

User FO files must be accessible from the database. They can be scattered in many directories on the server side. Instead of specifying a USER FO directory, users need to set the **DB_UsrFO** keyword to 1 in **dmconfig.ini** to enable the use of USER file objects. USER FOs are disabled by default.

Users can get the file name and file size of a FO by using the built-in functions, filename() and filelen().

Customizing the System File Object Path

DBMaker generates a series of file object subdirectories for storing system file objects. These subdirectories are located in the file object directory, which is specified by the **DB_FoDir** keyword in the **dmconfig.ini** file. When a file object subdirectory is filled to a threshold value by file objects, a new subdirectory is created. The threshold value is specified by the **DB_FoSub** keyword in the **dmconfig.ini** file, and may have a value from 100 to 10,000.

File object names take the form ZZ*xxxxxx.ext* where *xxxxxx* is a six digit base-36 serial number, and *ext* is the file extension of the object. The file extension of the object depends on the SET EXTNAME command. Refer to *System File Object Extension Names* for more information.

Subdirectories follow a naming convention based on the name of the first file object in the subdirectory. It takes the form SUB*xxxxxx* where *xxxxxx* is the six-digit base 36 number of the first file object to be inserted into the directory.

Although an FO directory can be shared by more than one database to simplify FO management, it is not recommended because it becomes inconvenient when backing up a database. The file object path may be changed before database startup by modifying the configuration parameter, or during runtime.

SETTING THE FO PATH OFFLINE

Users should specify where to put system FOs by setting the value of **DB_FoDir** in **dmconfig.ini**. The value of **DB_FoDir** is the full path of an existing directory. DBMaker must own the write-permission on that directory.

➲ **Example 1**

To create system FOs in the **/disk1/usr/fo** directory, add the following line to the **dmconfig.ini** file:

```
DB_FoDir = /disk1/usr/fo
```

➲ **Example 2**

To set DB_USRFO = 1 and enable user objects:

```
DB_USRFO = 1          ; enable USER file objects
```

SETTING THE FO PATH ONLINE

DBMaker provides a system procedure for users to modify the system file object directory while the database is running. This operation changes the setting of the following 3 items to the new value:

- **Run-time FO directory** — after the change is made, all new system file objects will be stored in the new FO directory.
- **DB_FoDir** — the next time the database is restarted, it will use the new FO directory.
- **$DB_FODIR alias.** —the default FO alias, which corresponds to the setting of the **DB_FoDir** keyword in **dmconfig.ini**.

➲ **Example**

To change the FO directory to a new directory, e.g. **/home/DBMaker/mydb/fo**, execute the following command:

```
dmSQL> call SETSYSTEMOPTION('fodir', '/home/DBMaker/mydb/fo');
```

In addition to being able to change the file object directory, users may query the system to return the current settings for the FO directory path.

➲ **Example**

The following command returns the current FO directory setting:

```
dmSQL> call GETSYSTEMOPTION('fodir', ?);
OPTION_VALUE: /home/DBMaker/mydb/fo
```

Generating File Objects

Several steps are required to generate file objects within DBMaker. First a FILE type column must be created on a table. Either system or user file objects may be inserted into an FO type column. To create a FO column, set the column type to FILE when creating the table.

➲ **Example 1**

To create a table named "person" with a file object column called "photo":

```
dmSQL> CREATE TABLE person (name CHAR(10), photo FILE);
```

➲ **Example 2**

If the FO to be input is on the on the server, DBMaker will link the FO column to the existing file and generate a user FO. If the FO is on the client side, DBMaker will create a system FO by copying the file from the client side to the FO directory on the server side. To insert FO data:

```
dmSQL> INSERT INTO person VALUES ('cathy','/disk1/image/cathy.bmp')
    2>;                                          // stored as a USER FO
dmSQL> INSERT INTO person VALUES ('jeff',?);
dmSQL/Val> &jeff.gif;                            // stored as a SYSTEM FO
dmSQL/Val> END;
```

Example 3

There are three varieties of fetching methods for a FO column: content, file name, and file size. To fetch an FO file named **cathy.bmp**:

```
dmSQL> SELECT photo, FILENAME(photo), FILELEN(photo) FROM person ;
  photo          filename(photo)           filelen(photo)
=========    =========================     ==============
012034451    /disk1/image/cathy.bmp                 21100
349045821    /disk1/usr/fo/ZZ000000.FOB             12034
```

For more information about manipulation on FO columns, refer to the *"JDBA Tool User's Guide*" and the *"ODBC Programmer's Guide*".

System File Object Extension Names

Users can set the system file object extension name using the SET EXTNAME command.

Example 1

To set the system file object <*extension_name*>:

```
SET EXTNAME TO <extension_name>;
```

There are two types of <*extension_name*>:

- A character string not over 7 characters in length, such as 'bmp', 'avi', 'jpg', etc.
- Using the SOURCE option, the extension name will be set equal that of the client's source file.

Example 2

To use the SET EXTNAME command:

```
dmSQL> CREATE TABLE t1 (c1 INT, f1 FILE);
dmSQL> INSERT INTO t1 (c1, f1) VALUES (?, ?);
dmSQL/Val> 1, &readme.txt;            //extension name : '.FOB'
1 rows inserted1
dmSQL/Val> SET EXTNAME TO doc
dmSQL/Val> 2, &readme.txt;            //extension name : '.doc'
```

```
1 rows inserted
dmSQL/Val> SET EXTNAME TO SOURCE;
dmSQL/Val> 3, &readme.txt;             //extension name : '.txt'
dmSQL/Val> END;
dmSQL> SELECT FILENAME(f1) FROM t1;
        c1                                                filename(f1)
==================
=========================================================
        1                   /usr1/fo/ZZ000001.FOB
        2                   /usr1/fo/ZZ000002.doc
        3                   /usr1/fo/ZZ000003.txt
3 rows selected
```

Updating File Objects

To update the contents of a FO column, use the SQL UPDATE command. DBMaker will replace the FO column with a new file.

As with inserting FOs, an FO column may be updated for a new SYSTEM FO or linked to a USER FO.

➲ **Example 3**

To link the **photo** column to **/disk2/image/common.bmp**:

```
dmSQL> UPDATE person SET photo = '/disk2/image/common.bmp' WHERE name =
'cathy';
```

Alternatively, a user can input new data from a file on the client side. For more information, refer to the "*JDBA Tool User's Guide*" 'and the "*ODBC Programmer's Guide*".

If the results of the UPDATE operation contain more than one tuple, only one file will be created. This file will be shared among the tuples to save disk space. DBMaker will maintain an internal counter to record how many tuples reference the file. In addition, if a user modifies the contents of the file through an external application program all tuples will recognize the modification.

When no tuples retain links to a system FO after UPDATE or DELETE operations, DBMaker will automatically delete the file after that transaction is committed. However, DBMaker never removes any USER FO even though there are no tuples referencing it, since DBMaker did not generate the file.

Renaming File Objects

Sometimes users need to change the positions or names of FOs because of full disks or a reorganization of the disk layout. DBMaker permits users to use the MOVE FILE OBJECT statement to change the name or path of the FO. Before using the MOVE FILE OBJECT command, use the operating system to move the files to the new location; DBMaker will make sure the new files exist before allowing the move.

➲ **Example 1**

To get the names of the files that will be moved using filename():

```
dmSQL> MOVE FILE OBJECT '/disk1/usr/fo/ZZ000000.FOB' TO
'/disk3/pub/photo1.bmp';
```

DBMaker also permits users to move FOs from one directory to another. Note that DBMaker permits using only one * character for the specified file name in the source directory, but does not allow using any * character in the destination directory. DBMaker does not support recursively-moving files. To move all of the files, not including subdirectories, from one directory to another, specify the former directory by adding the '/' or '/*' characters at the end of the directory.

➲ **Example 2**

Let there be four files in **/disk1/usr/fo** named **ABC1.FOB**, **ABC2.FOB**, **ABC3.FOB**, **ABC4.FOB.** To move **ABC1.FOB**, **ABC2.FOB**, **ABC3.FOB**, **ABC4.FOB** from **/disk1/usr/fo** to **/disk3/pub** file objects use:

```
dmSQL> MOVE FILE OBJECT '/disk1/usr/fo/ ' TO '/disk3/pub/ ';
dmSQL> MOVE FILE OBJECT '/disk1/usr/fo/* ' TO '/disk3/pub/ ';
dmSQL> MOVE FILE OBJECT '/disk1/usr/fo/*.FOB ' TO '/disk3/pub/ ';
dmSQL> MOVE FILE OBJECT '/disk1/usr/fo/A* ' TO '/disk3/pub/ ';
```

To move **ABC1.FOB** from **/disk1/usr/fo** to **/disk3/pub**:

```
dmSQL> MOVE FILE OBJECT '/disk1/usr/fo/*1.FOB ' TO '/disk3/pub/ ';
dmSQL> MOVE FILE OBJECT '/disk1/usr/fo/A*1.FOB ' TO '/disk3/pub/ ';
```

Predicate Operations on File Objects

As with BLOBs, users can test FOs for NULL values and use the CONTAIN and MATCH functions to perform pattern searches. Furthermore, the FO item can be used in arithmetic expressions with the FILELEN() built-in function, in Boolean expressions with the FILEEXIST () built-in function, and in string expressions with the FILENAME() built-in function.

In the event that a file has been removed or renamed from within the operating system, use the FILEEXIST() built-in-function to test which files exist. A value of 0 indicates that the referenced FO file does not exist, 1 indicates that it still exists, and NULL indicates that the tuple is a NULL value.

➲ **Example 1**

To select tuples with the .gif extension from the column **photo**:

```
dmSQL> SELECT * FROM person WHERE FILENAME(photo) LIKE '%.gif';
```

➲ **Example 2**

To fetch tuples with a size greater than 100KB from the column **photo**:

```
dmSQL> SELECT * FROM person WHERE FILELEN(photo) > 102400;
```

➲ **Example 3**

To fetch all tuples for files that exist:

```
dmSQL> SELECT * FROM person WHERE FILEEXIST(photo) 1;
```

File Object UNC Names

Universal Naming Convention (UNC) filenames can be used for the file object path and directory names in Microsoft Windows environments. This makes it easy to specify the path and directory names when a DBMaker server is running on a Microsoft Windows platform. Directories on machines other than the machine hosting the server can also be specified.

➲ **Example 1**

To retrieve all system FOs created in the **\\NTMACHINE\E\FO** directory when used in the **dmconfig.ini** file:

```
DB_FoDir = \\NTMACHINE\E\FO
```

➲ **Example 2**

To show how file objects work with UNC names:

```
dmSQL> CREATE TABLE t1 (c1 INT, c2 FILE);
dmSQL> INSERT INTO t1 VALUES (?, ?);
dmSQL/Val> 1, '\\NTMACHINE\D\DB\memo1.txt';
1 rows inserted
dmSQL/Val> 2, &c:\temp\memo2.txt;
1 rows inserted
dmSQL/Val> END;

dmSQL> SELECT c1, FILENAME(c2) FROM t1;
   c1                     FILENAME(c2)
========= ===================================================
    1           \\NTMACHINE\D\DB\memo1.txt
    2           \\NTMACHINE\E\FO\ZZ000001.txt


2 rows selected
```

File Object Path Default Aliases

DBMaker 3.73 supports two alias names for the file object path: $**DB_DbDir** and $**DB_FoDir**. The file object path alias provides an alias name to represent the real file object path. Users may insert/update/delete file objects through the alias path name. File objects can be moved more easily to another directory path.

The two alias names are set using the keywords **DB_DbDir** and **DB_FoDir**. in the **dmConfig.ini** file for $**DB_DbDir** and $**DB_FoDir**, respectively.

Example 1

The file object path alias is set to the path specified by the **DB_FoDir** keyword in **dmconfig.ini**:

```
…
DB_FoDir = "/usr1/tmp/employeedata/FO"
```

Example 2

To use the file object path alias to insert values into the FILE type column **photo**:

```
dmSQL> create table t1 (c1 INT ,photo FILE);
dmSQL> insert into t1 values (2, '$DB_FoDir/photo471.jpg')
```

In the above example, file **photo471.jpg** could also have been inserted using the full file object path '**/usr1/tmp/employeedata/FO/photo471.jpg**'.

FO and Applications

FILE type data is only supported by DBMaker, and is not defined by ODBC. Development tools, such as Inprise/Borland Delphi or Microsoft Visual Basic will not recognize FILE type data as valid. The configuration parameter **DB_FoTyp** may be used to specify which type of data FILE type data will be mapped to. To allow these tools to access data of FILE type, DBMaker should be set to internally map FILE type data to LONG VARBINARY by setting **DB_FoTyp** to 1. There is no mapping if **DB_FoTyp** is 0, and tools will fail to recognize FILE type data.

Example

To set the database to map FILE type data to LONG VARBINARY:

```
DB_FoTyp = 1
```

7.3 Journal of Large Objects

Logging transactions involving BLOB (LONG VARCHAR or LONG VARBINARY) data requires large amounts of disk space and results in decreased performance.

DBMaker lets the database administrator decide whether a BLOB in a specified tablespace is logged or not. DBMaker does not support logging for file objects (FILE).

DBMaker does not log the content of BLOB data by default. During the period between starting and shutting down the database, DBMaker ensures the consistency of BLOB data. Even in the event of a system crash, BLOB data is consistent after recovery.

However, when restoring a database from incremental backups, DBMaker does not ensure BLOB data consistency. Two steps must be taken to ensure that BLOB data is recorded in the journal. First, the configuration parameter **DB_BMode** must be set to record BLOB transactions. Second, the tablespace containing the BLOB data that is to be backed up must have been created with the BACKUP BLOB ON option.

BLOB Journal Logging

There are two conditions to ensure BLOB logging:

- Setting the value of the **DB_BMode** keyword in the ***dmconfig.ini*** file to 2 (BACKUP DATA AND BLOB mode).
- BLOB files are added to a tablespace that was created with the option BACKUP BLOB ON.

SETTING THE DB_BMODE VALUE

The keyword **DB_BMode** specifies the backup mode of a database. Setting the value to 0 enables NON-BACKUP mode, 1 enables BACKUP-DATA mode, and 2 enables BACKUP-DATA-AND-BLOB mode.

- NON-BACKUP (0) mode — cannot support incremental backup for tablespaces, including system or user-defined.
- BACKUP-DATA (1) mode — supports incremental backup for the system tablespace, data in user-defined tablespaces, but not BLOBs in user-defined tablespaces.

- BACKUP-DATA (2) mode — supports incremental backup for the system tablespace, data in user-defined tablespaces, BLOBs in user-defined tablespaces created with the BACKUP BLOB ON option, but not BLOBs in user-defined tablespaces created with the BACKUP BLOB OFF option.

To turn on BLOB journal logging, add a line to the ***dmconfig.ini*** file:

DB_BMode = 2 ;log all data including BLOB

For details on database backup mode, refer to Chapter 14, *Database Backup, Recovery, and Restoration.*

Setting the CREATE TABLESPACE BACKUP Option

The backup mode for an individual tablespace is set when it is being created. The syntax for the CREATE TABLESPACE command follows:

```
CREATE [AUTOEXTEND] TABLESPACE tablespace name [backup mode]
DATAFILE [tsfile , tsfile, ...];
```

Where:

```
backup mode ::= BACKUP BLOB OFF | BACKUP BLOB ON
tsfile ::= file_name TYPE = DATA | file_name TYPE = BLOB
```

Users can place important BLOBs in tablespaces with the BACKUP BLOB ON flag. It is a good idea to place BLOBs that do not need to be backed up in tablespaces with the BACKUP BLOB OFF setting in order to improve the system performance. Tablespace creators determine the trade-off.

Data and BLOB files must be specified in the dmconfig.ini file before the tablespace is created. This may be accomplished through the "user files" page in the JConfiguration Tool. Refer to the *JConfiguration Tool Reference* for detailed instructions on creating data and BLOB files.

➲ **Example 1**

To create tablespace **ts1** with BACKUP BLOB OFF and **ts2** with BACKUP BLOB ON:

```
dmSQL> CREATE TABLESPACE ts1 BACKUP BLOB OFF
    2> DATAFILE f1 TYPE = DATA, f2 TYPE = BLOB;
```

```
dmSQL> CREATE TABLESPACE ts2 BACKUP BLOB ON
    2> DATAFILE f3 TYPE = DATA, f4 TYPE = BLOB;
```

➲ **Example 2**

Query the **BK_MODE** column from the **SYSTABLESPACE** table to know the backup mode for each tablespace. The value 1 means that BACKUP BLOB is OFF, while 2 means that it is ON. Querying the backup mode of a tablespace will yield the following result:

```
dmSQL> SELECT TS NAME, BK MODE FROM SYSTABLESPACE;
     TS NAME        BK MODE
==============  ===========
SYSTABLESPACE         2
DEFTABLESPACE         2
ts1                     1
ts2                     2
4 rows selected
```

A summary of the interaction of the backup modes between a database and its tablespaces follows. 'Yes' indicates that the type of tablespace in question is backed up. 'No' indicates that it is not.

Database Backup Mode	**Tablespace Backup Mode**	**User-defined Tablespace (Data)**	**User-defined Tablespace (BLOB)**	**System Tablespace (Data and BLOB)**
NON BACKUP (**DB_BMode** = 0)	---	No	No	No
BACKUP DATA (**DB_BMode** = 1)	---	Yes	No	Yes
BACKUP DATA AND	BACKUP BLOB OFF	Yes	No	Yes

DATABASE BACKUP MODE	TABLESPACE BACKUP MODE	USER-DEFINED TABLESPACE (DATA)	USER-DEFINED TABLESPACE (BLOB)	SYSTEM TABLESPACE (DATA AND BLOB)
BLOB (**DB_BMode** = 2)	BACKUP BLOB ON	Yes	Yes	Yes

Before setting the backup mode, ensure that the journal file is large enough to record all BLOB data; otherwise, a journal full message may be returned. For information on how to adjust journal file size, refer to the subsection *Resizing Journal Space* in Chapter 5.

For concepts on data files, the BLOB file and the tablespace, refer to Chapter 5, *Storage Architecture*.

For information on the CREATE TABLESPACE command, refer to the *SQL Command and Function Reference.*

For information on the SYSTABLESPACE table, refer to Appendix B, *System Catalog Reference*.

File Object Journal Logging

DBMaker does not support journal logging of FOs. When backing up the database, back up all FOs belonging to the database as well by manually copying them into a backup directory. Alternatively, Bakup server may be used to automatically back up file objects to a backup directory. For more information on backing up file objects, refer to section 14.6, *Backup Server*. To determine what files belong to a database, query the SYSFILEOBJ table.

➲ **Example**

To retrieve the filenames of all FOs by querying the SYSFILEOBJ table:

```
dmSQL> SELECT FILE_NAME FROM SYSFILEOBJ;
```

Copy all FOs to the backup storage location. When restoring the database from a backup, copy all FOs as well. If the file paths or file names have changed, use the MOVE FILE OBJECT command to update the file names in the SYSFILEOBJ table.

7.4 Large Objects and SELECT INTO Command

The SELECT INTO command takes selected data and inserts it into a specified table. File objects and BLOBs can be moved from one table to another using this command. The SELECT INTO command can be used in a distributed database (DDB) environment.

In a local-to-local SELECT INTO statement, DBMaker needs to duplicate the BLOB data or increase the *shared counter* of the system file object or a user file object.

In a DDB environment, DBMaker copies the BLOB data from one site to another, but there are many considerations for file objects. DBMaker provides the distributed file object duplication mode (SET DFO DUPMODE command) to take care of processing file objects in a DDB environment.

SET DFO DUPMODE

The DFO DUPMODE tells a database whether file objects are to be copied to the target database or not. There are two modes for DFO DUPMODE: NULL and COPY mode.

➲ **Example**

Syntax for DFO DUPMODE: NULL and COPY mode:

```
dmSQL> SET DFO DUPMODE NULL;
dmSQL> SET DFO DUPMODE COPY;
```

SET DFO DUPMODE NULL

There are two cases in DDB mode to consider:

- The source and target databases are the same, including the local database or both remote databases. Since they are the same database, DBMaker only increases the shared counters of the file objects.
- The source and the target database are not the same. The target FILE column is set to NULL. Thus, the file objects in the source database are not sent out.

SET DFO DUPMODE COPY

There are three situations to consider for file objects.

- For user file objects, DBMaker only passes the source file name to the target database. The user needs to copy the files to a place where the target database can access them. Sometimes a user must use the UPDATE command or the MOVE FILE OBJECT command to change the file names on the target database if the new directories are not the same on the source database.
- For system file objects between two different databases, DBMaker will create a new system file object on the target database and copy the contents of the source database to it.
- For system file objects on the same database, local-to-local or remote-to-remote, DBMaker only increments the shared counters.

Limitations

DFO DUPMODE mode does not affect a SELECT INTO command used on a BLOB (LONG VARCHAR and LONG VARBINARY) column. BLOB data can be copied using the SELECT INTO command regardless of the DFO DUPMODE.

In a DDB environment, if a SELECT INTO command is used on a user file object and the option of DFO DUPMODE is set to COPY, then the user should be aware of the location of the linked file on the target database. The linked file object should exist in the same relative path on the target database. If it is not, the user should use the operating system to copy the file from the source to the target database and use the UPDATE or the MOVE FILE OBJECT commands for these columns if the file paths of the source and target databases are different.

If the user has not performed the above operations, an error message will be returned when querying the file object, because the file does not exist in the full path or the path of the file is incorrect.

When selecting a system file object from a remote database into the local database, DBMaker has to keep a record of the shared information. The information is kept within one SELECT INTO command. Therefore, there is still a duplicate file problem, which wastes space. Additionally, selecting system file objects into a remote database creates duplicate files.

The SET EXTNAME option does not affect the result of the SELECT INTO command. The extension names of the file objects on the source and target databases are the same. For example, the file name of the source database is 'ZZ000001.BMP', and then the target name of the file object on the target may be 'ZZXXXXXX.BMP'.

➲ **Example**

DBMaker assumes the data of CHAR, VARCHAR or BINARY as the file name, so users must make sure **db2** can access the **/etc/hosts** file from the view of **db1**. Select the CHAR column into the FILE column, where column **c2** in table **t2** on database **db2** is FILE type:

```
dmSQL> SELECT c1, '/etc/hosts' FROM db1:t1 INTO db2:t2(c1, c2);
```

Considering the FILE type column on the target database, the table below summarizes the effect of the SELECT INTO command with the different source data types:

TYPE ON THE SOURCE DATABASE	ENVIRONMENT	SET DFO DUPMODE	RESULT

Type on the Source Database	Environment	SET DFO DUPMODE	Result
string expression CHAR VARCHAR BINARY	non-DDB or DDB Environment.	…	Source: passes the file name. Target: inserts new user file objects.
FILE	The source and the target are the same database.	…	Increases the shared counter of the file objects.
	The source and the target database are not the same.	NULL	Target: inserts the NULL value.
		COPY	The source is the user file object: Source: passes the file name. Target: inserts the new user file object. The source is the system file object: Source: passes the content of the file object. Target: inserts the new system file object.
LONG VARCHAR LONG VARBINARY Other	…	…	Not supported.

8 Security Management

This chapter provides guidelines on setting up the security policies for a database, and includes information on security, authority levels, and table privileges.

8.1 Security Policies

DBMaker provides two kinds of security:

- **Database authority** — determines who can log on to DBMaker and the actions they can perform.
- **Object privileges** — controls access rights for DBMaker objects. DBMaker objects include tables, columns, views, domains, and synonyms.

8.2 Database Authority

Database authority is used to determine access for a database. DBMaker controls database access with *user names* and *passwords* and has four classes of users as shown in Table 12-1.

The SYSADM is the most powerful authority level in DBMaker. There can be only one SYSADM for every database. A user with SYSADM authority can grant DBA, RESOURCE or CONNECT authority to other users, and has all the privileges of the DBA authority level on the database.

Users with DBA authority level have all privileges for all objects in the database and can grant, change, or revoke object privileges for any user except users with SYSADM or DBA authority. They can also create new resources like tablespaces and files, and perform database administrative operations like starting/terminating and backing up databases.

Users with RESOURCE authority are allowed to create new tables or views, and to grant privileges on their own tables to other users.

Users with only CONNECT authority can access objects that they have been granted privileges for, but cannot create new tables or views. They may also select information from the system tables.

Authority levels are hierarchical

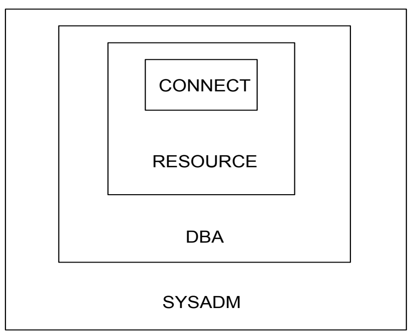

Figure 8-1: DBMaker database authority level hierarchy.

LEVEL	PRIVILEGES
SYSADM	Can grant and revoke security authority levels to all users except the SYSADM authority level. Can change the passwords of all users. Has all privileges of the DBA authority level.

LEVEL	PRIVILEGES
DBA	Has all privileges on tables except SYSTEM tables. Can grant/change/revoke object privileges of all users and groups. Can add/remove users from groups. Has privileges on database administration commands such as starting or terminating a database, creating/dropping/ altering a tablespace, and backing up a database. Has all the privileges of the CONNECT and RESOURCE authority levels.
RESOURCE	Can create and drop tables, views, domains, and synonyms. Can only drop tables, views, domains, and synonyms created by the user. Can grant/revoke owned table/view privileges to other users. Has any table privileges granted to the user. Has all the privileges of the CONNECT authority level.
CONNECT	Can log on to the database. Can select the SYSTEM tables. Has any table privileges granted to the user. This authority level must be granted before the other authority levels.

Table 8-1: DBMaker database authority levels

Managing Users

DBMaker provides several SQL commands for managing users. These commands allow new users to be added, existing users to be removed from a database, user passwords to be set or changed, and user authority levels to be granted.

ADDING A USER

The SYSADM must assign each user a user name and a password by using the GRANT (database authority) command before a user can log on.

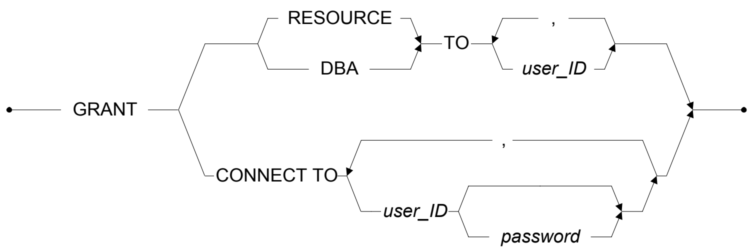

Figure 8-2 Syntax for the GRANT command

The GRANT command grants authority levels to users. Only the SYSADM can grant authority levels to other users. The SYSADM authority level cannot be granted to other users. As a result, for each database there is only one user with the SYSADM user name and SYSADM authority level. The SYSADM is also the default user who creates the database. Only the password can be changed for the SYSADM user name.

The SYSADM can grant CONNECT, RESOURCE or DBA authority to other users. If the GRANT command is used to grant RESOURCE or DBA authority to a user, it will not take effect until the next time the user connects.

The SYSADM can grant CONNECT authority to a user with a password. If the SYSADM does not specify the password, it means that user does not need a password to log on to database. A password can be any valid SQL identifier, which is not longer than eight bytes.

➲ **Example 1**

To grant CONNECT authority level and the password *jeff123* to user Jeff:

```
dmSQL> GRANT CONNECT TO Jeff jeff123;
```

➲ **Example 2**

To increase the authority level for user Jeff to RESOURCE:

```
dmSQL> GRANT RESOURCE TO Jeff;
```

➲ **Example 3**

To increase the authority level for user Jeff to DBA:

```
dmSQL> GRANT DBA TO Jeff;
```

CHANGING A PASSWORD

The ALTER PASSWORD command can be used to change a user's password.

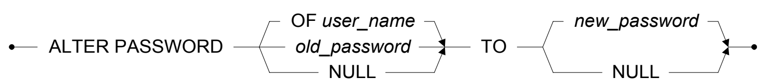

Figure 8-3 Syntax for the ALTER PASSWORD command

There are two ways to use the command:

- A user can change their own password with the ALTER PASSWORD <*old_password*> TO <*new_password*> command. The <*old_password*> must match the original password stored in the database.
- The SYSADM can change any user's password with the ALTER PASSWORD OF <*user_name*> TO <*new_password*> command. It is not necessary for the SYSADM to know the old password of other users.

➲ **Example 1**

The user Jeff changes his password from no password to *xyz@#*:

```
dmSQL> ALTER PASSWORD NULL TO “xyz@#”;
```

➲ **Example 2**

The SYSADM changes the password for user Jeff to ***xyz@#***:

```
dmSQL> ALTER PASSWORD OF Jeff TO “xyz@#”;
```

REMOVING A USER OR CHANGING A USER’S AUTHORITY LEVEL

The SQL REVOKE command removes a database authority level.

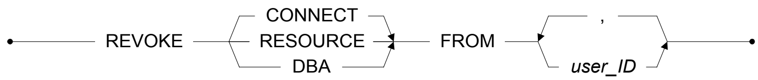

Figure 8-4 Syntax for REVOKE command

Revoking a user's RESOURCE or DBA authority does not take effect until the next time the user connects to the database.

Example 1

To revoke DBA authority from user Jeff:

```
dmSQL> REVOKE DBA FROM Jeff;
```

After executing the command, Jeff will no longer have DBA authority but will still have CONNECT authority.

Example 2

To remove Jeff's CONNECT authority and take away his ability to log on:

```
dmSQL> REVOKE CONNECT FROM Jeff;
```

Revoked Privilege	**Description**
DBA	Revoking DBA authority for a user means they can no longer create or drop tables and grant or revoke privileges from other users. The user will retain only the CONNECT authority unless granted the RESOURCE privilege. All tables, views, domains, and synonyms created by this user remain in the database.
RESOURCE	Revoking RESOURCE authority means the user can no longer create or drop tables. The user will retain only the CONNECT authority unless granted the DBA privilege. All tables, views, domains, and synonyms created by this user remain in the database.

REVOKED PRIVILEGE	DESCRIPTION
CONNECT	Revoking this authority means the user can no longer log on to the database. All privileges owned by this user on tables and views will be revoked. All tables, views, domains, and synonyms created by this user remain in database.

Table 8-3: Description of revoking DBMaker database authority levels

Managing Groups

To simplify management of authority levels, use a group to collect several users and/or other groups. Database privileges can then be granted to all members in a group at the same time with one command. Though a group is different from a user, it can be treated as a user. Object privileges granted to a group apply to all members in the group.

Only users with SYSADM or DBA authority levels can do the following:

- Create groups
- Add members to groups
- Remove members from groups
- Drop groups

CREATING GROUPS

The CREATE GROUP statement is used to create a new group.

```
•———— CREATE GROUP —— group_name ————•
```

Figure 8-5 Syntax for the CREATE GROUP command

The *group identification* (group name) uniquely identifies the name of a group in DBMaker. The group name cannot be SYSTEM, PUBLIC, GROUP or any existing user or group names.

➲ **Example**

To create a new group named COMMITTEE:

```
dmSQL> CREATE GROUP COMMITEE;
```

Adding Members to Groups

After creating a new group, users can be added using the ADD <*user name or group name*> TO GROUP command.

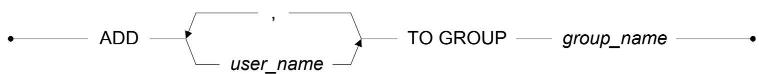

Figure 8-6 Syntax for the ADD … TO GROUP command

A group cannot be added as a new member of itself. Members of a group can include any existing user or group name.

➲ **Example**

To add user **Jeff** and group **RD** to the **COMMITEE** group and grant SELECT privilege to the **CASEMaker.EMPLOYEE** table:

```
dmSQL> ADD Jeff, RD TO GROUP COMMITEE;
dmSQL> GRANT SELECT ON CASEMaker.EMPLOYEE TO COMMITEE;
```

All members in **COMMITEE** will have the SELECT privilege for the **CASEMaker.EMPLOYEE** table.

Removing Members from Groups

The REMOVE <*user name or group name*> FROM GROUP command can be used to remove users from a specified group.

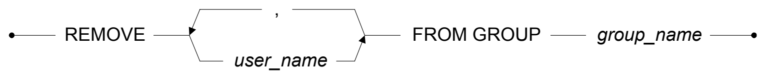

Figure 8-7 Syntax for the REMOVE … FROM GROUP command

The members removed from the group will lose all privileges granted to the specified group, but will retain privileges granted to them directly.

➲ **Example**

To remove user **Jeff** from the **COMMITEE** group:

```
dmSQL> REMOVE Jeff FROM GROUP COMMITEE;
```

After this command is executed, user **Jeff** will be removed from the group **COMMITEE** and lose SELECT privilege on the table **CASEMaker.EMPLOYEE**.

DROPPING GROUPS

The DROP GROUP command will drop a specified group from a database; all members in the group will lose the privileges granted for the group.

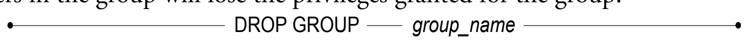

Figure 8-8 Syntax for the DROP GROUP command

➲ **Example**

To drop the **COMMITTEE** group from the database:

```
dmSQL> DROP GROUP COMMITEE;
```

8.3 Object Privileges

An *object* in a database includes the following items: tables, views, and columns in tables/views, domains, or synonyms. DBMaker provides security management for objects, which enables users to GRANT or REVOKE object privileges for other users.

All users can reference a domain by default, but only the creator can drop the domain. The privileges for a synonym are based on a base table. Refer to Chapter 6, *Managing Schema and* Schema Objects for detailed definitions of views, domains, and synonyms.

GRANTING OBJECT PRIVILEGES

The user that creates an object becomes the owner of the object and has all privileges for it. An owner can also grant privileges on the object to other users by using the SQL GRANT <*object privilege*> command.

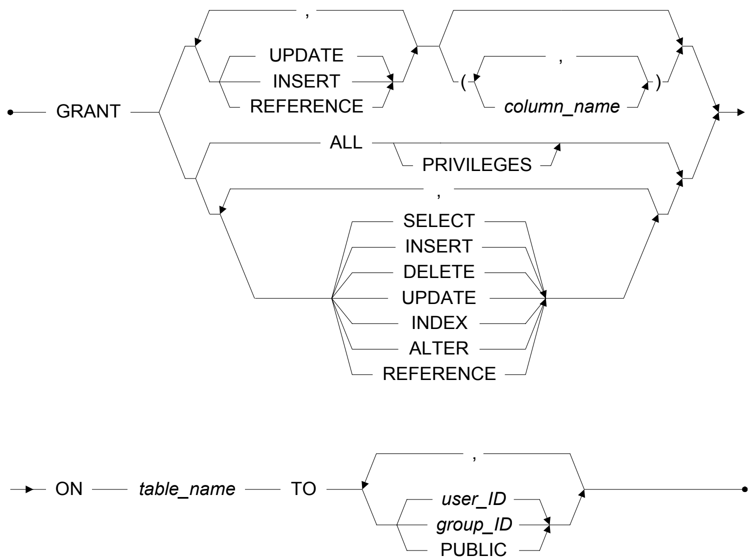

Figure 8-9 Syntax for the GRANT command

A user with DBA authority can grant privileges for any table or view in a database. A user with the RESOURCE authority can grant privileges only on tables or views created them. All privileges supported by DBMaker are described in Table 12-3.

INSERT, UPDATE, and DELETE privileges should be controlled to prevent corruption of information in a database. ALTER and INDEX privileges should be restricted to developers.

UPDATE, INSERT, and REFERENCE privileges can be restricted to some specific columns. Each column name must be qualified and be in *every* table identified in the ON clause.

PRIVILEGE	DESCRIPTION
SELECT	Allows users to select data from a table or view.
INSERT	Allows users to insert rows into a table or view and optionally insert into specified columns.
DELETE	Allows users to delete rows from a table or view.
UPDATE	Allows users to update a table or view and optionally update specified columns.
INDEX	Allows users to create or drop indexes for a table.
ALTER	Allows users to alter the definition of a table.
REFERENCE	Allows users to create a foreign key on a source table that references a primary key for a destination table or view.
ALL [PRIVILEGES]	Allows users to exercise all the above privileges for a table or view. PRIVILEGES is an optional keyword.

Table 8-5: Description for granting DBMaker table level privileges

The user in a GRANT command must have at least CONNECT authority. The group name is created using the CREATE GROUP command. The keyword PUBLIC includes all current and future users.

➲ **Example 1**

Jeff executes the GRANT command to give **Cathy** the read privilege to data in the **EMP_INFO** table, created by him:

```
dmSQL> GRANT SELECT ON EMP_INFO TO Cathy;
```

➲ **Example 2**

A DBA executes the GRANT command to give **Cathy** the read privilege to data in the **EMP_INFO** table created by **Jeff**:

```
dmSQL> GRANT SELECT ON Jeff.EMP_INFO TO Cathy;
```

➲ **Example 3**

A DBA gives INSERT and UPDATE privileges for the **PHONENO** column of **EMP_INFO** table to **Cathy**:

```
dmSQL> GRANT INSERT, UPDATE (PHONENO) ON Jeff.EMP_INFO TO Cathy;
```

Cathy will have no privileges for deleting information from the column.

➲ **Example 4**

Use of the PUBLIC keyword to permit all users to read data in the **Jeff.EMP_INFO** table:

```
dmSQL> GRANT SELECT ON Jeff.EMP_INFO TO PUBLIC;
```

Revoking Object Privileges

The REVOKE *<object privileges>* command revokes privileges granted to a user. The syntax for this command is shown in Figure 8-10.

The privileges in the REVOKE (object privileges) command are the same as those for the GRANT (object privileges) command. In the diagram, the user name represents an authorized user in the database, the group name represents a group of users, and the PUBLIC keyword represents all users in the database

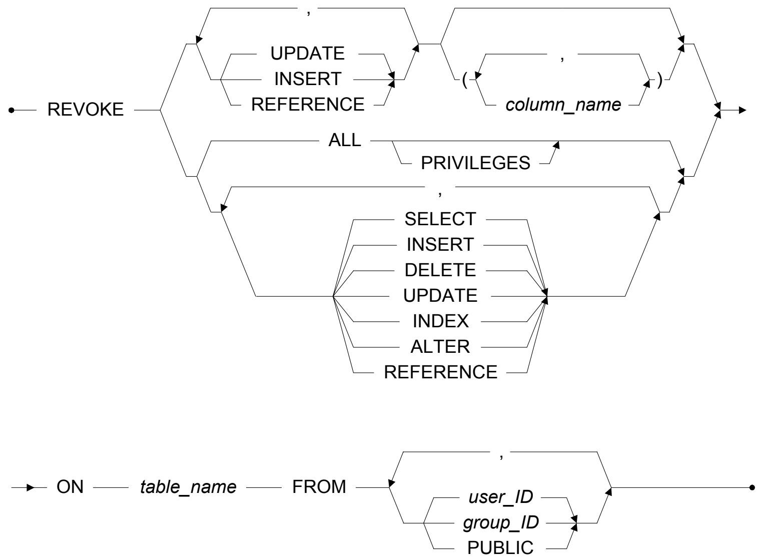

Figure 8-10 The REVOKE (object privileges) command

➲ **Example 1**

The following command revokes the SELECT privilege for the **EMP_INFO** table from **Cathy**:

```
dmSQL> REVOKE SELECT ON EMP_INFO FROM Cathy;
```

➲ **Example 2**

The following command revokes the SELECT privilege for table **Jeff.EMP_INFO** from **Cathy**:

```
dmSQL> REVOKE SELECT on Jeff.EMP_INFO FROM Cathy;
```

➲ **Example 3**

The following command revokes the UPDATE privileges on the column **PHONENO** in table **Jeff.EMP_INFO** from **group1**:

```
dmSQL> REVOKE UPDATE (PHONENO) on Jeff.EMP_INFO FROM group1;
```


➲ **Example 4**

The following command revokes all privileges granted to **PUBLIC** on the **EMP_INFO** table:

```
dmSQL> REVOKE ALL ON EMP_INFO FROM PUBLIC;
```

➲ **Example 5**

The following command revokes INSERT, UPDATE, and SELECT privileges for the **EMP_INFO** table from user **Cathy** and all users in **group2**:

```
dmSQL> REVOKE INSERT, UPDATE, SELECT ON EMP_INFO FROM Cathy, group2;
```

8.4 Security System Catalog

All information on authority levels, privileges, and groups is recorded in the following system catalogs:

- **SYSAUTHUSER**—authority level of each user.
- **SYSAUTHTABLE**—privileges on tables.
- **SYSAUTHCOL**—columns of a table to which a user has been restricted for INSERT, UPDATE, and REFERENCE privileges.
- **SYSAUTH**—group name, group creator, and number of group members.

The security system catalogs are owned by SYSTEM. No user (including SYSADM) can modify the system catalogs. See "Appendix B" for more details on the DBMaker system catalogs.

9 Concurrency Control

Transactions and concurrency control are described in this chapter. How DBMaker maintains concurrent access and data accuracy in a multi-user environment wit the lock mechanism is also described. Section 9.1 presents the transaction concept and the functions used in managing a transaction. Section 9.2 describes the necessity of concurrency control in a database system. Finally, section 9.3 explains concurrency control techniques used by DBMaker.

9.1 Transactions

In a database, a *transaction* is a work unit that is composed of one or more SQL statements. It is an *atomic* operation. That means it should either complete a series of statements entirely or do nothing at all. Serial, atomic, permanent, consistent, and isolated are the properties of a transaction.

Transaction States

A transaction must be in one of the following states:

- **Active—**When a transaction starts to execute, it immediately goes into an active state. In the active state, a transaction can perform various database operations.
- **Partially Committed—**When a transaction reaches its last statement in DBMaker (such as COMMIT WORK), it enters into the partially committed

state. The transaction has completed its execution and can still be aborted if an error occurs during the actual output. The result cannot be written to disk and a hardware failure may preclude its successful completion.

- **Committed—**When a transaction has completed its execution successfully it enters into the committed state.
- **Failed—**When a transaction cannot proceed to a normal conclusion, it enters into the failed state. This may be caused by hardware or logic errors, or a user abort of the transaction during an active state.
- **Aborted—**When a transaction has ended unsuccessfully, it enters into the aborted state. In this situation, any change or effect that a transaction has applied to the database must be rolled back.

The state diagram corresponding to a transaction is shown in Figure 9-1.

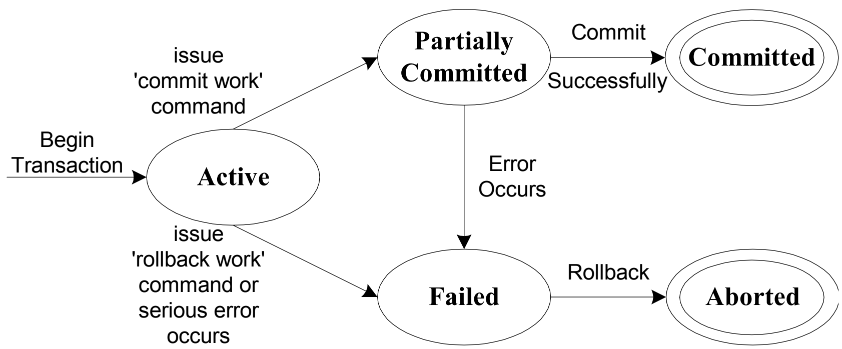

Figure 9-1: The transaction states

Managing a Transaction

When connecting to DBMaker, a transaction starts automatically and enters the active state. DBMaker will automatically begin a new transaction after the preceding transaction has been terminated.

Every time a statement executes a transaction is committed automatically by DBMaker. This is known as *autocommit mode*. In this mode, the lifetime of a transaction equals the lifetime of a single SQL statement. That means when one transaction is terminated at the end of an SQL statement, another begins with the next SQL statement. Each SQL statement is an independent transaction.

To force a transaction to remain uncommitted until several SQL statements have been executed, change to *manual commit mode* by issuing a SET AUTOCOMMIT OFF command. In this mode, a transaction can only be committed by using the SQL command COMMIT WORK. As many SQL statements as necessary can be executed before ending the transaction. To end the transaction, issue a COMMIT WORK command to commit changes, or issue a ROLLBACK WORK command to abort any changes made and terminate the transaction.

To return to autocommit mode, issue a SET AUTOCOMMIT ON command. The default transaction mode is AUTOCOMMIT ON.

NOTE *After a transaction is terminated, all resources allocated are released.*

Using a Savepoint

A *savepoint* is an intermediate point that can be arbitrarily declared within the context of a transaction. A savepoint is used to rollback the work performed after a savepoint has been declared within a transaction.

For example, a transaction with a series of statements is executed, and an error occurs while executing the twentieth statement. If a savepoint is marked between the fifteenth and sixteenth statements, the first fifteen statements can be preserved. A user can roll back to the savepoint and begin issuing commands from the sixteenth SQL statement after correcting the error. The user does not need to abort the transaction and resubmit all the statements. Figure 9-2 shows an example of this.

However, if the user does not mark a savepoint between the fifteenth and sixteenth statements, the transaction must be aborted and the first fifteen statements resubmitted. This is inconvenient and wastes time. A savepoint solves this problem completely.

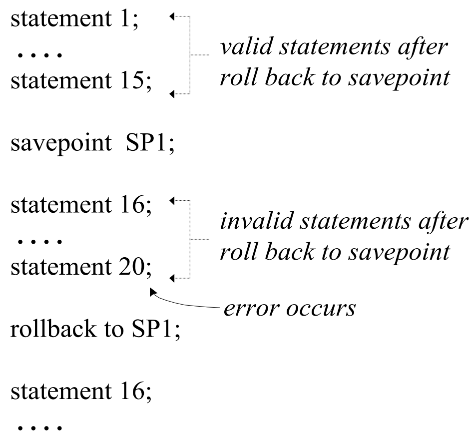

Figure 9-2: Using Savepoints

The SAVEPOINT and ROLLBACK TO … commands mark a savepoint and rollback to a specific savepoint.

➲ **Example 1**

The SAVEPOINT command:

```
dmSQL> SAVEPOINT <savepoint_name>;
```

➲ **Example 2**

The ROLLBACK TO … command:

```
dmSQL> ROLLBACK TO <savepoint_name>;
```

The user specifies the <*savepoint_name*>. After rolling back to a savepoint, the system resources that were allocated after the savepoint, like locks, are released.

9.2 Multi-User Environment

When more than one user is accessing a database, consider what can happen when they try to access data simultaneously.

Sessions

A *connection* is a communication pathway between a user and DBMaker. A communication pathway is established using shared memory or a network.

Before using the database resources, establish a connection to DBMaker using the following SQL statement.

➲ **Example**

To connect a user to a DBMaker database:

```
dmSQL> CONNECT TO database_name user_name password;
```

When a user connects to a DBMaker database, the specific connection is called a *session*. A session lasts from the time a user connects to a DBMaker database until the time the user disconnects from it. A session can only have one active transaction at a time.

The Necessity of Concurrency Control

In a multi-user database system environment, more than one user can connect to a database at the same time. This could possibly result in many transactions updating the same database simultaneously.

If no concurrency control mechanism is used, several situations could result in data inconsistency:

- The lost update problem.
- The temporary update problem.
- The incorrect summary problem.

Lost Update Problem

A lost update problem occurs when two transactions update a data item at approximately the same time.

➲ **Example**

Transactions T1 and T2 read and modify the value of X but use different calculations to modify the value. This results in the transactions each containing a different value for X. T1 writes the value it holds for X to the database after it is read but before it is written by T2. T2 then writes the value it holds for X to the database, overwriting the value written by T1. The value written by T1 is lost:

```
      T1                        T2
--------------            --------------
read(X);
                          read(X);
X = X - N;
                          X = X + M;
write(X);
                          write(X);
```

Temporary Update Problem

A temporary update problem occurs when a transaction updates a value, but is rolled back after another transaction updates the same value.

➲ **Example**

Transaction T1 reads and modifies the value of X, writes it back to the database, and then continues with other commands. While transaction T1 continues executing, transaction T2 reads the value of X, modifies it to a new value, and writes it back to the database. Transaction T1 then fails before completion, and must roll back all values to restore the database to its original status. The database management system restores the original value of X, overwriting the value written by transaction T2. The value of X calculated by transaction T2 exists only temporarily:

```
      T1                        T2
--------------            --------------
read(X);                      read(X);
```

```
X = X - N;
write(X);                       X = X + M;
                                 write(X);
rollback;
```

INCORRECT SUMMARY PROBLEM

An incorrect summary problem occurs when a transaction is calculating the aggregate sum of a number of records while other transactions are updating those records.

➲ **Example**

Transaction T1 calculates the aggregate sum using the values of X and Y at the same time transaction T2 is modifying those values. Transaction T2 updates the value of X before transaction T1 uses it to calculate the sum, and updates the value of Y after transaction T1 uses it to calculate the sum. This results in transaction T1 using some values to calculate the sum before they are updated, and using others after they are updated. When both transactions complete, the value of the sum is incorrect with respect to the values in the database:

```
      T1                        T2
--------------          --------------
sum = 0;
                        read(X);
                        X = X - N;
                        write(X);
read(X);
sum = sum + X;
read(Y);
sum = sum + Y;
                        read(Y);
                        Y = Y + N;
                        write(Y);
```

There are various techniques to solve concurrency problems, such as locks and time stamps. The next section shows how the locking technique is applied in DBMaker to control concurrent execution of transactions.

9.3 Locks

In this section, the lock concept is first presented. Then, the DBMaker lock mechanism is introduced, including lock granularity and lock modes. Finally, dealing with deadlock is demonstrated.

Lock Concept

In general, a multi-user database system uses several forms of locking to synchronize the access of concurrent transactions. Before accessing the data objects, such as tables and tuples, a transaction must lock those data objects.

DBMaker locking is fully automatic and does not require any user action. Implicit locking occurs in all SQL statements; the users do not need to explicitly lock any data objects in the database.

Shared and Exclusive Locks

In general, two types of locking are used to allow multiple-read with single-write operations in a multi-user database.

- **Share Locks** (S)—A transaction involving a read operation on a data object. To support a higher degree of data concurrency, several transactions can acquire share locks on the same data object at the same time.
- **Exclusive Locks** (X)—A transaction involving an update operation on a data object. This transaction is the only one that can access the object until the exclusive lock is released.

Two-Phase Locking

The two-phase locking protocol is used to ensure the transactions are serialized. In the two-phase locking protocol, each transaction must issue all lock requests before it can issue any unlock requests.

The protocol can be divided into two phases:

- **Expanding (growing) phase**—This phase allows the transaction to issue any new lock requests that are required. Unlock requests are not permitted in this phase.
- **Shrinking phase**—This phase allows the transaction to release locks acquired in the expanding phase. New lock requests are not permitted in this phase.

The two-phase locking protocol is currently used by DBMaker to provide concurrency control by serializing transactions.

DEADLOCK

When two or more transactions are waiting for the release of data locked by other transactions before it can proceed, a deadlock occurs.

➲ **Example**

T1 is waiting for T2 to release the share lock of X, while T2 is waiting for T1 to release the share lock of Y. Therefore, deadlock occurs and the system will wait indefinitely:

```
      T1                        T2
--------------          --------------
share lock(Y);
read(Y);
                        share lock(X);
                        read(X);
exclusive lock(X);
(T1 waits for T2)       exclusive lock(Y);
                        (T2 waits for T1)
```

Lock Granularity

There are three granularity levels for data locks in DBMaker: relation (table), page, and tuple (row). A relation contains several pages, and a page contains several tuples.

A lock applied on a higher level carries through to lower levels. For example, if a user gets an exclusive lock (X lock) on a relation, all pages and tuples that are included in this relation will have the X lock applied to them. Therefore, no user can access any tuple or page from this relation. However, if a user gets an X lock on a tuple, another user can get an X lock on another tuple simultaneously. There is no interference between two objects at the same level when using the X lock. Figure 9-3 shows the lock granularity (levels) in DBMaker.

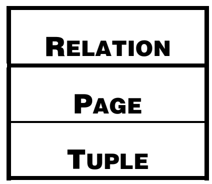

Figure 9-3: Lock granularity

Using a higher lock granularity results in a lower degree of data concurrency, in contrast, the higher lock granularity uses fewer system resources (such as shared memory). Selecting the lock granularity level is a trade-off between concurrency and resources. In DBMaker, the default lock granularity level is page, but if a different lock granularity is required, it can be specified when creating a table. Refer to Chapter 5, *Storage Architecture* for more information.

Lock Types

The main lock modes (types) supported in DBMaker are shared (S) and exclusive (X) locks. More than one user can have an S lock on a data object simultaneously, but only one user can have an X lock on a data object. In addition to S and X locks, another lock mode called an *intention lock* is supported.

When a data object is locked, the system will automatically assign an intention lock to the next higher granularity object. For example, an S lock specified on a tuple will generate an intention S (IS) lock on the page which includes this tuple, and an IS lock on the relation which the tuple belongs to.

The supported intention lock modes are:

- **IS**—Indicates that the S lock is specified at a lower granularity.
- **IX**—Indicates that the X lock is specified at a lower granularity.
- **SIX**—Indicates that an S lock is specified at the current granularity and an X lock is specified at a lower granularity. This is a combination of S and IX locks.

The result from the compatibility of each of the lock modes is listed in Table 9-1. T represents true, which means the matrix for each of the two lock modes are compatible and can exist on a data object simultaneously. F represents false, which means the matrix for each of the two lock modes are not compatible and cannot exist simultaneously.

If lock requests on a data object conflicts with an existing lock on that object, this request will not execute until the existing lock is released, or until the waiting time for the lock request times out. If the error message 'Lock timeout' is returned to the user, the waiting time for the lock has expired. The default waiting time is 5 seconds. However, users can specify a different waiting time by setting the value of the **DB_LTimO** keyword in the **dmconfig.ini** file to another value according to their individual requirements.

➲ **Example**

The following shows how to set the waiting time to 8 seconds:

```
DB_LTimO = 8;
```

	IS	**S**	**IX**	**SIX**	**X**
IS	T	T	T	T	F
S	T	T	F	F	F
IX	T	F	T	F	F
SIX	T	F	F	F	F
X	F	F	F	F	F

Table 9-1: Compatibility matrix for lock modes

Dealing with Deadlock

By analyzing the "wait for" graph, DBMaker can automatically detect a deadlock situation. If a deadlock is detected, a victim transaction will be aborted to solve the deadlock problem.

➲ **Example**

DBMaker detects a deadlock when transaction T2 issues an X lock on Y. Transaction T2 will be aborted to resolve the deadlock problem and the user executing transaction T2 will receive the error message, "transaction aborted due to deadlock**"**:

```
      T1                      T2
--------------          ------------
share lock(Y);
read(Y);
                        share lock(X);
                        read(X);
exclusive lock(X);
(T1 waits for T2)       exclusive lock(Y);
                        (T2 waits for T1)

                        T2 aborted by DBMaker
```

10 Triggers

Triggers are a very useful and powerful feature of the DBMaker database server. Triggers automatically execute predefined commands in response to specific events, regardless of which user or application program generated them.

Triggers allow a database to be customized in ways that may not be possible with standard SQL commands. The database can consistently control complex or unconventional database operations without requiring any action on the part of users or application programs.

Use triggers to:

- Implement business rules.
- Create an audit trail for database activities.
- Derive additional values from existing data.
- Replicate data across multiple tables.
- Perform security authorization procedures.
- Control data integrity.
- Define unconventional integrity constraints.

Exercise restraint when using triggers to avoid forming complex interdependencies within the database that may be difficult to follow and change. Use triggers only when the desired functionality cannot be implemented using standard SQL commands and integrity constraints.

10.1 Trigger Components

DBMaker stores trigger definitions in the system catalog.

Every DBMaker trigger has six main components:

- **Trigger Name**—a name that uniquely identifies the trigger
- **Trigger Action Time**—the time relative to when a trigger will be fired
- **Trigger Event**—a specific situation that occurs in the database in response to some user action, such as inserting data into a table
- **Trigger Table**—the name of the table the trigger executes on
- **Trigger Action**—an SQL statement or stored procedure that is executed when the trigger event occurs
- **Trigger Type**—the type of trigger

Each of these components must be present in all triggers. In addition, there is an optional component, the REFERENCING clause.

Trigger Name

The trigger name uniquely identifies a trigger. Trigger names have a maximum length of 18 characters and may contain letters, numbers, the underscore character, and the symbols # and $. The first character cannot contain a number, and the name cannot contain spaces.

Trigger Action Time

The trigger action time specifies whether it should fire before or after the SQL statement that activates it. The trigger action time is specified by the BEFORE and AFTER time keywords. The BEFORE keyword instructs the trigger to fire before the trigger statement. The AFTER keyword instructs the trigger to fire after the trigger statement. Only one trigger time can be specified for each trigger.

Trigger Event

The trigger event is the database operation that causes a trigger to operate, or *fire*. The trigger event may be an INSERT, UPDATE, or DELETE statement that operates on the trigger table. There can be only one trigger event for each trigger statement. However, multiple trigger events can be used to activate multiple triggers.

Trigger Table

The trigger event operates on the associated trigger table. The trigger table must be a base table; it cannot be a temporary table, view, or synonym. A trigger may only have one trigger table.

Trigger Action

A trigger action is the command that a trigger executes when it fires. The trigger action may be an INSERT, UPDATE, DELETE, or EXECUTE PROCEDURE statement. A trigger can only have a single trigger action.

Trigger Type

The trigger type specifies how many times the trigger will fire for each trigger event. There are two types of triggers: row triggers and statement triggers. The FOR EACH ROW option specifies a row trigger, which fires a trigger action once for each row modified by the trigger event. The FOR EACH STATEMENT option specifies a statement trigger, which fires a trigger action once for each trigger event.

REFERENCING Clause

The REFERENCING clause defines correlating names for the old and new values in a column. This is primarily used when the default OLD and NEW names cannot be used because of a name conflict with a table.

10.2 Trigger Operation

DBMaker checks to see if a trigger should be fired and will execute the defined triggers each time a user or an application program causes a trigger event. Firing triggers from within a database ensures that DBMaker handles data consistently across all applications. This guarantees that when a specific event occurs, a related action is also performed.

Users can create triggers to implement domain, column, referential, and unconventional integrity constraints. However, these can also be done by declarative integrity control.

Triggers do not have an owner, but are associated with a table.

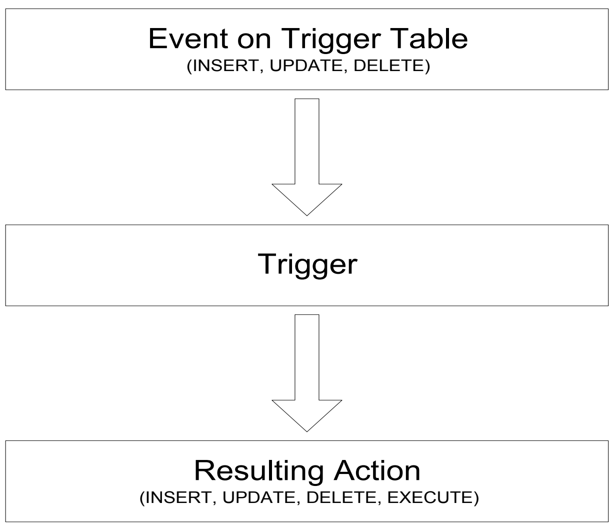

Figure 10-110-2: Trigger event and action

10.3 Creating Triggers

The CREATE TRIGGER command creates a new trigger associated with a specific table. Only a user with privilege on the trigger table can execute the command. The user must also have the necessary object privileges for all objects referenced in the trigger definition in order to successfully create a trigger.

Basic Requirements

All of the CREATE TRIGGER statements must contain at least the following:

- A trigger name
- The trigger action time (before or after)
- The trigger event
- The trigger table
- The trigger type (row or statement)
- The trigger action

Security Privileges

All SQL statements in the trigger action operate with the same privileges as the owner of the trigger table, and not with the privileges of the user executing the trigger event. If the trigger exists, any user executing the trigger event will result in the trigger firing.

CREATE TRIGGER Syntax

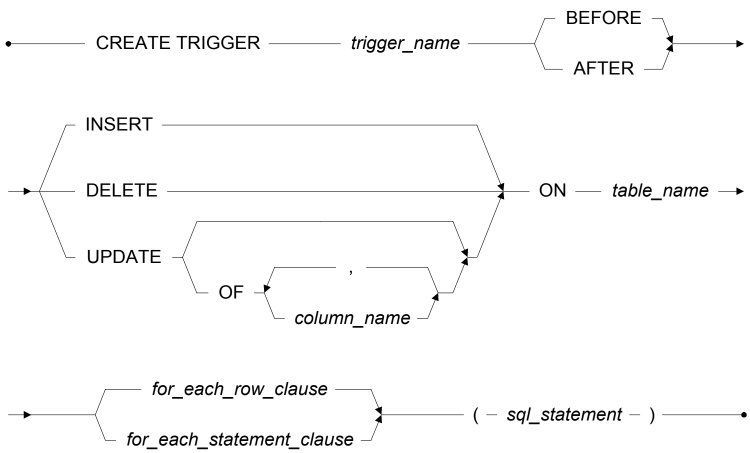

FOR EACH ROW clause

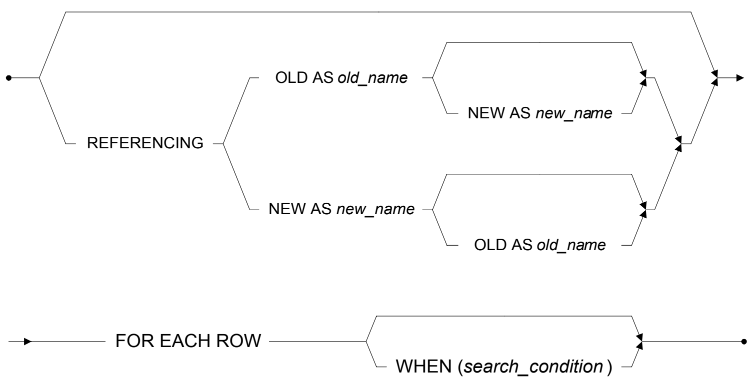

FOR EACH STATEMENT clause

FOR EACH STATEMENT

Figure 10-3: The Syntax for the CREATE TRIGGER Statement

Specifying the Trigger Action Time

You can use the trigger time and trigger type in combination to create four triggers for each table for the same event (INSERT, DELETE, or UPDATE). For each event the BEFORE/FOR EACH ROW, AFTER/FOR EACH ROW, BEFORE/FOR EACH STATEMENT, and AFTER/FOR EACH STATEMENT combinations are possible.

A BEFORE/FOR EACH STATEMENT trigger executes once and only once before the triggering statement is performed. That is before the occurrence of the trigger event. An AFTER/FOR EACH STATEMENT trigger executes once and only once after the triggering statement is complete. Note that BEFORE and AFTER statement triggers are executed even if the triggering statement does not process any rows.

BEFORE OR AFTER INSERT OR DELETE TRIGGER EVENTS

The following examples show how to create triggers that fire before or after INSERT or DELETE trigger events. The trigger action is represented by <*sql_statement*>.

➲ **Example 1**

To define four triggers for an INSERT event on table **t1**:

```
dmSQL> CREATE TRIGGER tr1 BEFORE INSERT ON t1 FOR EACH STATEMENT <sql_statement>
dmSQL> CREATE TRIGGER tr2 BEFORE INSERT ON t1 FOR EACH ROW <sql_statement>
dmSQL> CREATE TRIGGER tr3 AFTER INSERT ON t1 FOR EACH ROW <sql_statement>
dmSQL> CREATE TRIGGER tr4 AFTER INSERT ON t1 FOR EACH STATEMENT <sql_statement>
```

➲ **Example 2**

To define four triggers for a DELETE event on table **t1**:

```
dmSQL> CREATE TRIGGER tr1 BEFORE DELETE ON t1 FOR EACH STATEMENT <sql_statement>
dmSQL> CREATE TRIGGER tr1 BEFORE DELETE ON t1 FOR EACH ROW <sql _statement>
dmSQL> CREATE TRIGGER tr1 AFTER DELETE ON t1 FOR EACH ROW <sql _statement>
```

```
dmSQL> CREATE TRIGGER tr1 AFTER DELETE ON t1 FOR EACH STATEMENT <sql_statement>
```

BEFORE OR AFTER THE UPDATE TRIGGER EVENT

The situation is different for UPDATE events. Two types of UPDATE triggers can be created: UPDATE <*table*> triggers, or UPDATE OF <*column*> triggers. An UPDATE <*table*> trigger fires whenever the table is updated. An UPDATE OF <*column*> trigger fires when specific columns are updated. Either one UPDATE <*table*> trigger or multiple UPDATE OF <*column*> triggers can be created on a single table. UPDATE OF <*column*> triggers may contain multiple columns, but columns in all UPDATE OF <*column*> triggers in a table must be mutually exclusive.

➲ **Example 1**

To create a column trigger **tr1** on table **tb1** of columns **c1**, **c2**, that has four columns, **c1**, **c2**, **c3**, and **c4**:

```
dmSQL> CREATE TRIGGER tr1 AFTER UPDATE OF c1,c2 ON tb1
                          FOR EACH ROW
                          (INSERT INTO tb2 VALUES (old.c1, old.c2));
```

If a second UPDATE column trigger **tr2** that specifies column **c2** is created, the command will fail because **c2** already appears in trigger **tr1**:

```
dmSQL> CREATE TRIGGER tr2 AFTER UPDATE OF c2,c3 ON tb1
                          FOR EACH ROW
                          (INSERT INTO tb3 VALUES (old.c2, old.c3));
ERROR (6150): [DBMaker] the insert/update value type is incompatible with column
data type or compare/operand value is incompatible with column data type in
expression/predicate
```

If there are four columns in a table, you can create at most four UPDATE column triggers or one UPDATE table trigger, for triggers of the same type (for instance, a BEFORE/FOR EACH ROW trigger).

FOR EACH ROW / FOR EACH STATEMENT Clause

The FOR EACH STATEMENT clause specifies that a trigger will fire once and only once for each trigger event. Even if the trigger event statement does not process any rows, the trigger will fire.

The FOR EACH ROW clause specifies that a trigger will fire once for each row that the trigger event modifies. If the trigger event does not modify any rows, the trigger will not fire. The OLD and NEW keywords are used to identify which values from the trigger table are to be used in the trigger action. The OLD keyword indicates that trigger table values from before the trigger event are used in the trigger action. The NEW keyword indicates that trigger table values from after the trigger event are used in the trigger action.

➲ **Example 1**

The following statement shows how to create an UPDATE column trigger on table **Sales**. The **totSales** field is a calculated field from other two fields **unitPrice** and **unitSale**. Both **unitPrice** and **unitSale** are triggering columns.

```
dmSQL> CREATE TRIGGER trTotalSale AFTER UPDATE OF unitPrice, unitSale ON Sales
                                  FOR EACH ROW
                                  (UPDATE Sales
                                  SET totSales = new.unitPrice * new.unitSale);
```

➲ **Example 2**

In this example, there are four triggers.

```
dmSQL> CREATE TRIGGER trig1 BEFORE UPDATE ON Orders
                     FOR EACH STATEMENT
                    (EXECUTE PROCEDURE checkPrivilege);

dmSQL> CREATE TRIGGER trig2 BEFORE UPDATE ON Orders
                     FOR EACH ROW
                    (INSERT INTO Log_Old_Value (old.customer, old.amount));

dmSQL> CREATE TRIGGER trig3 AFTER UPDATE ON Orders
                     FOR EACH ROW
                    (INSERT INTO Log_New_Value (new.customer, new.amount));

dmSQL> CREATE TRIGGER trig4 AFTER UPDATE ON Orders
                     FOR EACH STATEMENT
                    (EXECUTE PROCEDURE Log_Time);
```

If a user executes an UPDATE statement that changes two rows of the **Orders** table, the effect and order of the execution is as follows:

1. Procedure checkPrivilege is called.
2. Insert one row to Log_Old_Value table.
3. Update one row.
4. Insert one row to Log_New_Value table.
5. Procedure Log_Time is called.
6. Insert one row to Log_Old_Value table.
7. Update one row.
8. Insert one row to Log_New_Value table.
9. Procedure Log_Time is called.

Stored procedures cannot contain COMMIT, ROLLBACK, or SAVEPOINT transaction control statements. Triggers can specify only a single triggered action, which must be enclosed in parentheses.

Using the Referencing Clause

In row triggers, the <*sql_statement*> (or action body) should indicate whether the column values used are from before or after the trigger event. For example, to log the old price and new price when updating the price of a sale item, use the keywords OLD and NEW as shown in example 2 in the section “FOR EACH ROW / FOR EACH STATEMENT Clause”.

However, in some rare cases the tables may contain columns with the names NEW or OLD. If this is the case, use the referencing clause to define correlation names. The reference clause allows for the creation of two prefixes that can be used with a column name: one to reference the old value of the column, and one to reference the new value. These prefixes are called correlation names. Use the keywords OLD and NEW to indicate the correlation names.

➲ **Example**

```
dmSQL> CREATE TRIGGER tr_log_price AFTER UPDATE OF price ON New
                                   REFERENCING OLD as pre NEW as post
                                   FOR EACH ROW
                                   (INSERT INTO logTbl
                                    VALUES (item_no, today(), pre.price,
                                    post.price));
```

In this example, the triggering table name is NEW, so the correlation names **pre** and **post** are used in the action body. Referencing clauses are only valid for row triggers, and are not allowed in statement triggers.

If a trigger event is INSERT, there is no old value for the newly inserted record, so the old value is not available. Similarly, if the trigger event is DELETE, there is no new value for the deleted record, so the new value is not available. For an UPDATE event trigger, both old and new values are available.

Using the WHEN Condition

A WHEN condition clause may precede a FOR EACH ROW triggered action to make the action execution dependent on the result of a boolean expression. The WHEN clause consists of a keyword WHEN followed by the conditional statement in parentheses. The WHEN clause follows the action time and precedes the triggered action body. The WHEN clause is not allowed in the definition of a statement trigger, it is only allowed in row trigger.

➲ **Example 1**

The following trigger will log a customer complaint into the **logComplain** table when a customer calls to complain about something (Assume the call code '**c**' means it is a complaint call.).

```
dmSQL> CREATE TRIGGER tr_log_complain INSERT ON Customer_Call
                                      FOR EACH ROW
                                      WHEN (new.call_code = 'c')
                                      (INSERT INTO logComplain
                                       VALUES (Today(), Cus_Name));
```

If the WHEN condition is included in a trigger definition, the WHEN clause is evaluated for each row. If the WHEN condition evaluates to TRUE for a row, the triggered action is fired for that row. If the WHEN condition evaluates to FALSE or unknown for a row, the triggered action is not fired for that row.

The result of WHEN condition only affects the execution of the triggered action, it has no effect on the triggering statement.

Example 2

To create three triggers to record all INSERT, UPDATE and DELETE operations on table **emp**:

```
dmSQL> CREATE TRIGGER trig_emp_insert AFTER INSERT ON emp
                                      FOR EACH ROW
                                      (INSERT INTO emp_audit
                                       VALUES (NULL, NULL,
                                       new.empId, new.empName));

dmSQL> CREATE TRIGGER trig_emp_update AFTER UPDATE ON emp
                                      FOR EACH ROW
                                      (INSERT INTO emp_audit
                                       VALUES (old.empId, old.empName,
                                       new.empId, new.empName));

dmSQL> CREATE TRIGGER trig_emp_update AFTER DELETE ON emp
                                      FOR EACH ROW
                                      (INSERT INTO emp_audit
                                       VALUES (old.empId, old.empName,
                                       NULL, NULL));
```

Example 3

If a primary key is changed, all the foreign keys can be changed in cascade. Suppose **deptNo** is the primary key on table dept, **DeptNo** is foreign key on table **emp**.

```
dmSQL> CREATE TRIGGER trig_upd_dept BEFORE UPDATE OF deptNo ON dept
                                    FOR EACH ROW
                                    WHEN (NEW.deptNo <> OLD.deptNo)
                                    (UPDATE emp SET emp.DeptNo = NEW.deptNo
                                     WHERE emp.DeptNo = OLD.deptNo);
```

Example 4

If the primary key is deleted, all the foreign keys can be deleted in cascade.

```
dmSQL> CREATE TRIGGER trig_del_dept BEFORE DELETE ON dept
                                    FOR EACH ROW
                                    (DELETE FROM emp
                                     WHERE emp.DeptNo = OLD.deptNo);
```

➲ **Example 5**

If a primary key is updated, all the foreign keys can be set to NULL.

```
dmSQL> CREATE TRIGGER trig_del_dept BEFORE UPDATE ON dept
                                    FOR EACH ROW
                                    (UPDATE emp set DeptNo = NULL
                                     WHERE emp.DeptNo = OLD.deptNo);
```

➲ **Example 6**

If the number of parts in stock is lower than a given level, the parts should be reordered. The part number and quantity will be recorded to a table called **pending_orders** for further action.

Inventory: part_no int, parts_on_hand int, reorder_level int, reorder_qty int

pending_orders: part_no int, qty int, order_date date

```
dmSQL> CREATE TRIGGER tr_reorder AFTER UPDATE OF parts_on_hand ON Inventory
                                 FOR EACH ROW
                                 WHEN (new.parts_on_hand < new.reorder_level)
                                 (INSERT INTO pending_orders
                                  VALUES (new.part_no, new.reorder_qty, today()))
```

Specifying the Trigger Action

The trigger action is the SQL statement that is performed when the trigger event occurs. The trigger action can be an INSERT, DELETE, UPDATE, or EXECUTE PROCEDURE statement. No other statements are allowed. Stored procedures cannot contain COMMIT, ROLLBACK or SAVEPOINT transaction control statements. Triggers can specify only a single trigger action, which must be enclosed in parentheses.

➲ **Example**

The following statement creates a trigger on table **emp**.

```
dmSQL> CREATE TRIGGER trigExample AFTER INSERT ON emp
                                  FOR EACH ROW WHEN (new.empNo > 0)
                                  (INSERT INTO personnel(new.empName,
                                   new.empAddress, new.Manager));
```

In this example, the trigger name is **trigExample**. The AFTER option is specified, which means this trigger will be fired after the INSERT statement executes on table **emp**. The triggering event is INSERT, the triggering table is **emp**. The trigger type is FOR EACH ROW. The SQL action that is triggered is INSERT.

```
dmSQL> CREATE TRIGGER trDelAcct AFTER DELETE ON Account
                                FOR EACH ROW
                                (INSERT INTO oldAccount
                                 VALUES (Old.Customer_name));
```

In the above example, the trigger **trDelAcct** will add the deleted customer name into the **oldAccount** table when one deletes a record from the **Account** table. You cannot create a trigger on a temporary table, view, or system table.

10.4 Modifying a Trigger

A trigger cannot be modified, but its definition can be replaced. When you want to modify a trigger definition, use the ALTER TRIGGER statement.

ALTER TRIGGER Syntax

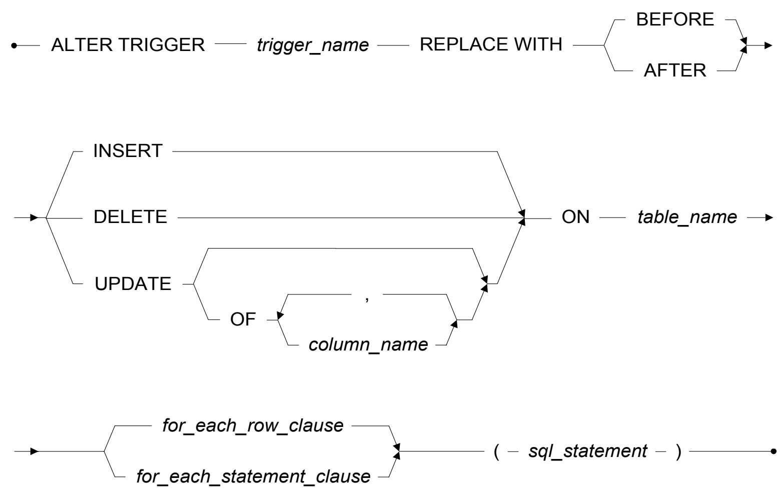

FOR EACH ROW clause

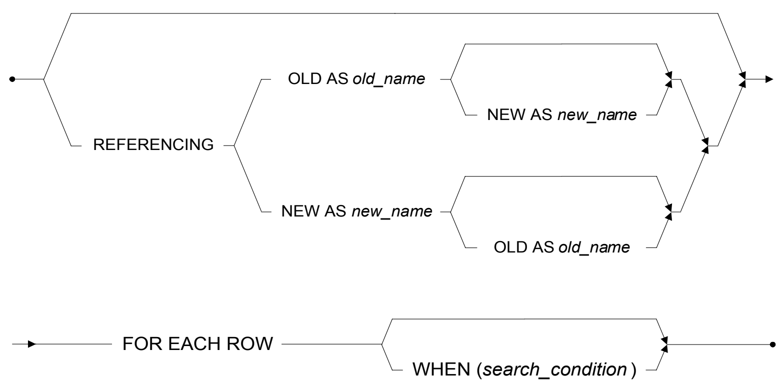

FOR EACH STATEMENT clause

Figure 10-4 Syntax for the ALTER TRIGGER command

Replacing a Trigger Action

To replace a trigger action, use the statement ALTER TRIGGER tr1 REPLACE WITH ...

➲ **Example 1**

If a manager quits then their data needs to be deleted from the **manager** table. To create a trigger on the **employee** table:

```
dmSQL> CREATE TRIGGER delEmp AFTER DELETE ON employee
                            FOR EACH ROW
                            ( DELETE FROM manager WHERE empId = old.empId )
```

➲ **Example 2**

It is possible to add another condition to the triggered action, such as "delete the data from manager table only when the employee is a project manager". To replace a trigger action on the **employee** table and add a condition:

```
dmSQL> ALTER TRIGGER delEmp REPLACE WITH AFTER DELETE ON employee
                           FOR EACH ROW
                           ( DELETE FROM manager
                             WHERE empId = old.empId
                             AND title = 'Project Mananger')
```

Alternatively, the trigger can be dropped and recreated.

10.5 Dropping a Trigger

The DROP TRIGGER statement can be used to delete a trigger from the database.

DROP TRIGGER Syntax

•—— DROP TRIGGER —— *trigger_name* —— FROM —— *table_name* ——•

Dropping the Trigger

Deleting a table will cause triggers referencing the table to be deleted. When a table schema is altered, DBMaker will try to execute the trigger according to the new table definition the next time the trigger is executed. If the specified column in a triggering event or action is dropped, the trigger execution and statement will fail. The only solution is to drop the trigger or modify the trigger definition according to the new table schema. To drop a trigger, specify the name of the trigger to delete, and the associated table.

➲ **Example 1**

To drop the **myTrigger** trigger from **myTable** table:

```
dmSQL> DROP TRIGGER myTrigger FROM myTable;
```

➲ **Example 2**

To create a trigger named **tr1** for table **t1**:

```
dmSQL> CREATE TRIGGER tr1 AFTER UPDATE ON t1
                          FOR EACH ROW
                          ( DELETE FROM t2 WHERE c1 = old.c1 )
```

If the column **c1** in table **t2** is dropped or the type is changed, an execution error will occur when the triggering statement (update on **t1**) is performed causing the DBMS to attempt to fire trigger **tr1**.

10.6 Using Triggers

There are several ways to use triggers.

Stored Procedures in Action Body

One of the most powerful features of a trigger is the ability to use a stored procedure as a trigger action. The EXECUTE PROCEDURE statement calls a stored procedure, enabling you to pass data from the triggering table to the stored procedure and then execute the procedure.

➲ **Example**

To create a trigger and use the EXECUTE PROCEDURE statement:

```
dmSQL> CREATE TRIGGER trLogPrice AFTER UPDATE OF price ON Sales
                                FOR EACH ROW
                                (EXECUTE PROCEDURE
                                logPrice(item_no, new.price, old.price));
```

Users can pass values to a stored procedure in the argument list. If the stored procedure call is part of the action for a row trigger, users can use the OLD and NEW correlation values to pass the column values it. If the stored procedure is part of an action statement trigger, users can only pass constants to the stored procedure.

Within a trigger action, you can update non-triggering columns in the triggering table, with or without a stored procedure. A stored procedure fired by a trigger cannot contain transaction control statements, like BEGIN WORK, COMMIT WORK, ROLLBACK WORK, SAVEPOINT, or DDL statements.

The stored procedure as a trigger action cannot be a cursory procedure that returns more than one row.

Trigger Execution Order

The column numbers in the triggering columns determine the order of trigger execution. The trigger execution begins with the trigger with the smallest triggering column number and proceeds in order to the highest number. In the following example a=column1, b=column 2, c=column 3 and d= column 4.

➲ **Example**

The operation UPDATE **t1** SET b=b+1, c=c+1 will fire both triggers. Trigger **trig1**, having a lower triggering column number than **trig2**, will be executed first. The following assumes four columns named **a**, **b**, **c**, and **d** from table **t1**.

```
dmSQL> CREATE TRIGGER trig1 AFTER UPDATE OF a,c ON t1
                            FOR EACH STATEMENT (UPDATE t2 set c1=c1+1)

dmSQL> CREATE TRIGGER trig2 AFTER UPDATE OF b,d ON t1
                            FOR EACH STATEMENT (UPDATE t2 set c2=c2+1)
```

Security and Triggers

First, the user must have permission to run the trigger event; otherwise, the user cannot trigger the event. However, the user does not have to have permission to run the triggered action because the SQL statements in the triggered action operate under the domain privilege of the trigger owner. Once a trigger is created successfully, the trigger creator has privilege to execute the triggered action. Any one else who can issue the triggering statement can also fire the trigger.

➲ **Example**

User **B** can update on both tables **T1** and **T2**, and user **A** can update **T1**, but not **T2**. Now user **B** creates a trigger on update **T1**, and the action updates **T2**. When user **A** updates **T1**, the triggered action (**update T2**) is executed successfully since the triggered action is running under the domain privilege of user **B**. This security rule simplifies execution and eliminates the requirement for the user to have more privileges to execute the triggered action.

Cursors and Triggers

UPDATE or DELETE statements within a cursor act differently than a single update or delete statement. The entire trigger will be executed with each update or delete with the WHERE CURRENT OF clause.

For example, if four rows are changed with a cursor, the BEFORE/FOR EACH STATEMENT, BEFORE/FOR EACH ROW, AFTER/FOR EACH STATEMENT and AFTER/FOR EACH ROW triggers will be executed four times - once for each row.

Cascading Triggers

Executing one trigger may cause another trigger to also be executed. You can use cascading triggers to enforce referential integrity. DBMaker supports a maximum of 64 cascading triggers.

➲ **Example 1**

To first delete a customer from the customer table, trigger the action to delete customer related records in the order table, which in turn will trigger the action to delete order related records in the item table:

```
dmSQL> CREATE TRIGGER cas1 AFTER DELETE ON customer
                          FOR EACH ROW
                          (DELETE FROM orders WHERE cust_num = old.cust_num);

dmSQL> CREATE TRIGGER cas2 AFTER DELETE ON orders
                          FOR EACH ROW
                          (DELETE FROM items WHERE order_num = old.order_num);
```

In DBMaker, if users create recursive triggers, it will not return an error at trigger creation time. However, users will get an error when the recursive triggers execute and meet the maximum limit of cascading trigger levels.

10.7 Enabling and Disabling Triggers

When a trigger is created, the trigger is in **enabled** mode, which means the triggered action executes when the trigger event occurs.

Sometimes users may need to disable a trigger:

- When users have to load a large amount of data, disabling the triggers temporarily will speed up the loading operation.
- When the objects referenced in a trigger are unavailable.

➲ **Example 1**

To disable trigger tr1 for table t1:

```
dmSQL> ALTER TRIGGER tr1 ON t1 DISABLE
```

➲ **Example 2**

To enable trigger tr1 on table t1:

```
dmSQL> ALTER TRIGGER tr1 ON t1 ENABLE
```

In summary, a trigger has two possible modes:

- **Enabled**—A trigger is enabled when created. The triggered action fires when the event occurs.
- **Disabled**—A disabled trigger does not execute, even if the event occurs.

10.8 Create Trigger Privileges

To create a trigger for a table, a user must be the table owner or DBA. The trigger creator must have privileges to all objects referenced in the CREATE TRIGGER statement to be successful.

In DBMaker, a trigger has no owner; it is associated with a table. The table owner and DBA have all privileges associated with a trigger. They can create, drop, or alter the triggers.

The SQL statements in the trigger action operate under the domain privileges of the trigger owner, instead of the domain privileges of the user executing the trigger event.

11 Stored Commands

A stored command is a compiled SQL DML statement stored in the database. A stored command is precompiled in an executable format. The same command can be executed without compiling and optimizing. It is possible to create a stored command for any frequently used SQL statement, achieving better performance. Stored commands cab be considered a subset of stored procedures that only contain one SQL statement without program logic.

11.1 Creating Stored Commands

Use the CREATE COMMAND statement to create a stored command.

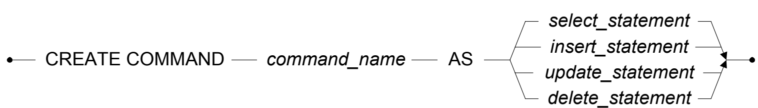

Figure 11-1 Syntax for the CREATE COMMAND ststement

Input parameters in the SQL statement can be used when creating a stored command. The actual value of the input parameters for a stored command can be assigned at the time of execution.

➲ **Example 1**

To create a stored command named **sc1** for the SQL DML statement using a table with the definition **t1** (**c1** INT, **c2** INT, **c3** CHAR(32)):

```
dmSQL> INSERT INTO t1 VALUES (1, ?, ?)
```

Example 2

Alternatively use:

```
dmSQL> CREATE COMMAND sc1 AS INSERT INTO t1 VALUES (1, ?, ?)
```

Example 3

To create stored commands for other DML statements:

```
dmSQL> CREATE COMMAND sc2 AS SELECT c1, c2 FROM t1
dmSQL> CREATE COMMAND sc3 AS UPDATE t1 SET c1 = c2+1 WHERE c2 > ?
dmSQL> CREATE COMMAND sc4 AS DELETE FROM t1 WHERE c2 > ?
```

After creating a stored command, a user with permission can execute it directly using dmSQL or in an application program. If the stored command has input parameters, its value can be determined using parameter marks, constants, NULL, DEFAULT, or built-in functions (built-in functions can't have arguments), when executing the stored command. When a user executes a stored command, the number of input parameters should be equal to the number of input parameters in the stored command.

11.2 Executing a Stored Command

Use the EXECUTE COMMAND statement to execute a stored command.

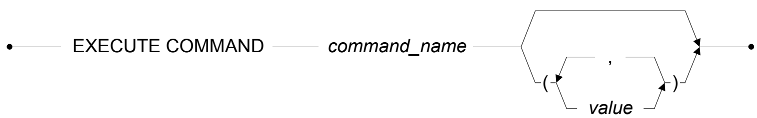

Figure 11-2 Syntax for the EXECUTE COMMAND statement

Example 1

```
dmSQL> EXECUTE COMMAND sc1 (200, 'john')
```

Example 2

```
dmSQL> EXECUTE COMMAND sc1 (DEFAULT, ?)
```

Example 3

```
dmSQL> EXECUTE COMMAND sc1 (?, NULL)
```

Example 4

```
dmSQL> EXECUTE COMMAND sc1 (?, ?)
```

A stored command may be dropped when it is no longer useful.

11.3 Dropping a Stored Command

Use the DROP COMMAND statement to drop a stored command.

Figure 11-3 The syntax of DROP COMMAND statement

➲ **Example**

```
dmSQL> DROP COMMAND sc1
```

11.4 Stored Command Security

Stored commands are treated similarly to other database schema objects, and as a result, any user must consider security and object privileges when creating or using one.

Only the creator or users that have the RESOURCE privilege can create a stored command. A user can only create a stored command from an SQL DML statement if the user has privileges to execute the SQL DML statement.

➲ **Example**

User **joe** with resource privilege wants to create a stored command **CheckDate** with the following syntax:

```
dmSQL> CREATE COMMAND CheckDate AS SELECT FirstName, LastName, Hiredate FROM
SYSADM.Employee WHERE HireDate > '1995-01-01';
```

The employee table is owned by SYSADM, so the system administrator must first grant select permission to user **joe** on the table **SYSADM.employee** before user **joe** can create the stored command.

A user must have the execute privilege for a stored command to execute it. In order to allow a stored command to be used by others, the user can grant the execute privilege on a stored command. However, only users with the necessary privileges (DBA, SYSADM, the creator of the stored command, or others granted the privilege) may grant or revoke execute privileges for stored commands.

A DBA has the execute privilege on all stored commands in a database. The owner of a stored command has execute, grant, and revoke privileges. Only the owner of stored command can drop it.

Granting Execute Privilege

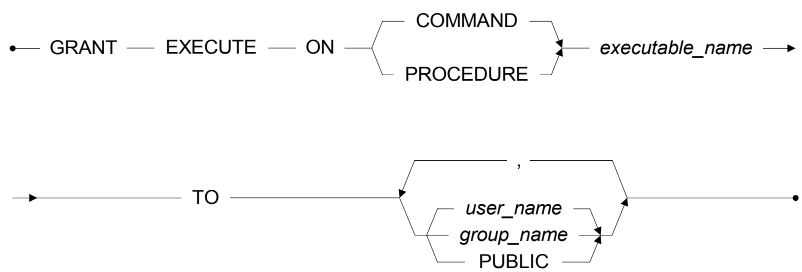

Figure 11-4 Syntax for the GRANT EXECUTE privilege

➲ **Example**

To grant **usr1** the EXECUTE privileges for commands on **sc1**:

```
dmSQL> GRANT EXECUTE ON COMMAND sc1 TO usr1
```

Revoking Execute Privileges

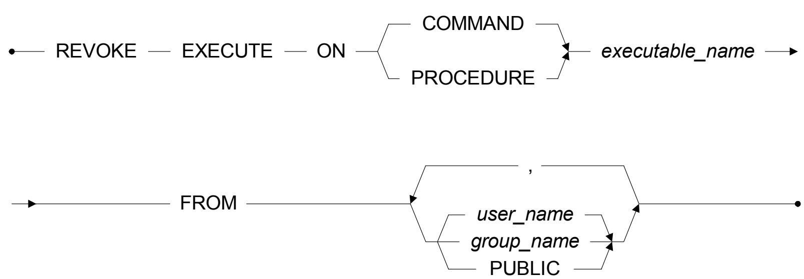

Figure 11-5 Syntax for the REVOKE EXECUTE privilege

➲ **Example**

To revoke the EXECUTE privileges from **usr1** for **sc1**:

```
dmSQL> REVOKE EXECUTE ON COMMAND sc1 FROM usr1
```

11.5 Lifecycle of a Stored Command

A stored command will be invalid if one of the related tables in the stored command is dropped or altered. If any programs were written previously using old column information, it may cause unpredictable results at time of execution.

The benefit of stored command is improved performance when repeatedly executing a SQL command. DBMaker also considers when the execution plan should be refreshed, such as UPDATE STATISTICS. When the UPDATE STATISTICS command is issued, all execution plans for the stored command will be updated to achieve better performance.

11.6 Getting Information for Stored Commands

Users can get information about stored commands from the system table SYSCMDINFO. Table 11-1 lists the columns of the SYSCMDINFO table and their values.

COLUMN NAME	VALUE	COMMENT
MODULENAME	Module name that the stored command belongs to	This column is used by an ESQL application or stored procedure. It can be ignored if it is a pure stored command.
CMDNAME	Stored command name	none
CMDOWNER	Stored command owner	none
STATEMENT	Original SQL string	none
NUM_PARM	Number of parameters	none
STATUS	0, 1, or 2	0 - invalid stored command. It cannot be executed.

COLUMN NAME	VALUE	COMMENT
		1 – valid stored command. It can be executed 2 – the stored command needs to be rebound. It can be executed after internal rebinding.

Table 11-1: Details of the SYSCMDINFO table

Users can get stored command information by issuing the following statement in dmSQL:

```
dmSQL> SELECT * from SYSCMDINFO;
```

12 Stored Procedures

A stored procedure is a special kind of user-defined function that contains embedded SQL statements. Once the stored procedure has been created, it is stored in executable format in the database as an object. This allows the database engine to bypass repeated SQL compilation and optimization, increasing the performance of frequently repeated tasks. The stored procedure is executed as a command in interactive SQL, or invoked in application programs, trigger actions, or other stored procedures.

Accomplish a wide range of objectives with stored procedures including improving database performance, simplifying the writing of applications, and limiting or monitoring access to a database.

Because a stored procedure is stored as an executable object in the database, it is available to every application running on the database. Several applications can use the same stored procedure to reduce development time for an application.

12.1 Creating Stored Procedures

A stored procedure is an ESQL/C program. Stored procedures can perform any function a C application can, including calling other C functions and system calls. Therefore, a C compiler is required for writing stored procedures.

An ESQL/C program for a stored procedure consists of a CREATE PROCEDURE statement, a declare section if needed, and the code section. If your program does not use any host variables, the declare section can be omitted.

Example

To create a stored procedure called **a_phone** with one input parameter, one output parameter, and a return value (**status**):

```
EXEC SQL CREATE PROCEDURE a_phone (CHAR(13) name, CHAR(13) phone OUTPUT)
                RETURNS STATUS;
{
    EXEC SQL BEGIN CODE SECTION;

        EXEC SQL SELECT PHONE FROM TBL WHERE NAME = :name INTO :phone;

        EXEC SQL RETURNS STATUS SQLCODE;

    EXEC SQL END CODE SECTION;
}
```

The structure of this program will be explained in the following sections.

Create Procedure Syntax

In the head of a procedure definition is a CREATE PROCEDURE statement. The syntax for the CREATE PROCEDURE statement is:

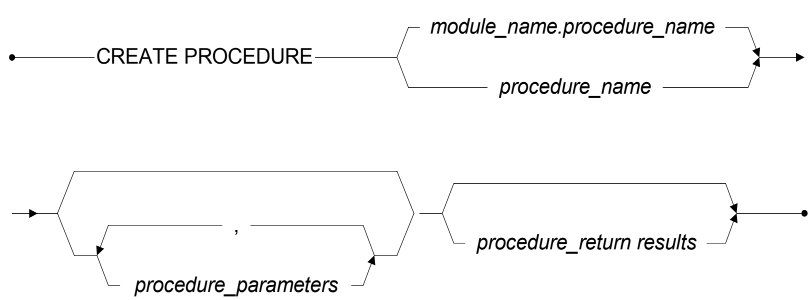

<*procedure_parameters*> clause

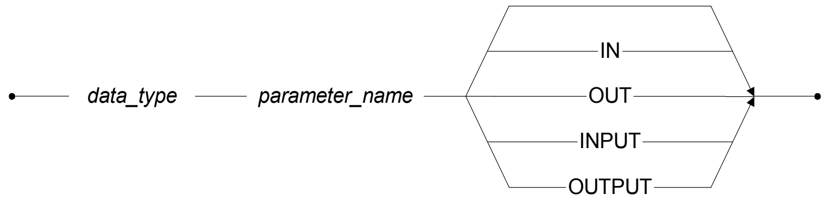

<*procedure_return_result* > clause

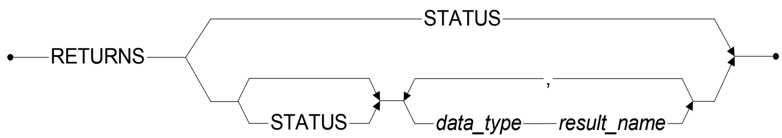

Figure 12-1 Syntax for the CREATE PROCEDURE statement

➲ **Example 1;**

The following are examples of the syntax of the CREATE PROCEDURE statement.

```
dmSQL> CREATE PROCEDURE p1 (INTEGER n IN) RETURNS STATUS;
dmSQL> CREATE PROCEDURE p2 (INTEGER n1 IN, INTEGER n2 OUTPUT) RETURNS CHAR(12)
nm;
dmSQL> CREATE PROCEDURE p3 (CHAR(10) par1 OUTPUT, SMALLINT par1)
       RETURNS STATUS, TIMESTAMP ret1, FLOAT ret2;
```

In a CREATE PROCEDURE statement the procedure name and the name and type of any I/O parameters must be provided.

Using Parameters

If parameters are required, a list of type-name pairs for the parameters must be given in parentheses. IN/OUT (or INPUT/OUTPUT) parameter attributes must be put after each type-name pairs. If there is no parameter attribute, IN will be used by default. Input parameters are used to pass a value to a procedure. In example 1, there is one input parameter **name**. Every time the procedure is executed, a value for the input parameter must be provided.

Output parameters are used to get a single result, not a result set, after the procedure is executed. In example 1, procedure **a_phone** has an output parameter, **phone**. The output parameter must have a buffer assigned to it to receive the result. After the procedure executes, the phone number for the name inputted can be retrieved from the buffer.

The result list is required for a stored procedure to retrieve a result set of tuples from the database. If the procedure does not return selected results then there is no need for the result list. The keyword RETURNS is used to start the result list. It is a list of name type pairs.

The STATUS keyword is used to indicate an integer value to be returned after the procedure executes.

➲ **Example**

To execute a procedure with one input parameter and a value:

```
EXEC SQL CREATE PROCEDURE t19 (FLOAT ifl) RETURNS STATUS,
                                                  FLOAT fl,
                                                 DOUBLE db;
{
    EXEC SQL BEGIN CODE SECTION;
    EXEC SQL RETURNS STATUS SQLCODE;
    EXEC SQL RETURNS SELECT fl, db FROM t8 WHERE fl < :ifl into :fl, :db;
    EXEC SQL END CODE SECTION;
}
```

DBMaker now supports the following data types for input and output parameters: INTEGER, SMALLINT, SERIAL, CHAR(), DATE, TIME, TIMESTAMP, FLOAT, DOUBLE, REAL.

Return Select Statement

A procedure can return a result set using the host variable mechanism to pass information to the user executing the stored procedure. In the code of the stored procedure use the RETURNS keyword to instruct the preprocessor to generate a host variable related to C code. The RETURNS keyword precedes the select statement that produces the result set.

➲ **Example**

There are two RETURNS in this example, one in the create procedure statement and the other in the select statement forming a pair. If there is to be a result set returned, declare output parameters with RETURNS in the create procedure statement and put the RETURNS keyword in the select statement:

```
EXEC SQL CREATE PROCEDURE  get_all_phone RETURNS CHAR(12) name, CHAR(12) phone;
{
    EXEC SQL BEGIN CODE SECTION;
        EXEC SQL RETURNS SELECT NAME, PHONE FROM TBL INTO :name, :phone;
    EXEC SQL END CODE SECTION;
}
```

Module Names

When a user creates a stored procedure DBMaker will use the owner name and procedure name as the default dynamic link library name. The user can call or drop his or her stored procedures using only the procedure name. Any user can call another user's procedure using the full procedure name: *owner.procedure_name*.

A user can also specify a module name in the CREATE PROCEDURE syntax to change the default dynamic link library name. If a module name is specified in the CREATE PROCEDURE syntax, users will need to call or drop the procedure with the full procedure name; *module_name.owner.procedure_name*, even if the user created it.

Variable Declaration

The host variables in stored procedures are declared in the same way as in ESQL/C. The declare section in a stored procedure must be put before the code section, not in ESQL/C programs. C variables can be placed before or after the declare section, but need to be before the code section.

Code Section

All statements should be in the code section except the variable declaration. Any non-declaration statement before the code section may cause problems resulting in compile errors or wrong results being returned. Statements after the CODE SECTION will not be executed.

Configuration Settings for Stored Procedures

When a stored procedure is created, a corresponding dynamic link library is built and stored on the server. By default, the library file is placed in the DBMaker server's working directory. The database administrator can set a preferred path to place the library files for stored procedures using the configuration keyword **DB_SPDir**.

The keyword **DB_SPLog** is used by client users to set the directory they prefer to receive error message files and trace log files, transmitted from the database server while creating or executing stored procedures.

➲ **Example 1**

To set the default path of dynamic link library files for stored procedures to **/usr1/dbmaker/data/SP** add the following line in the **dmconfig.ini** file:

```
DB_SPDIR=/usr1/dbmaker/data/SP
```

➲ **Example 2**

To set the stored procedure log file directory to **c:\usr\jerry\data\SP** add the following line in the **dmconfig.ini** file:

```
DB_SPLOG=c:\usr\jerry\data\SP
```

Creating a New Stored Procedure from File

First, write the stored procedure and save it to a file, then use DBMaker tools like dmSQL or DBATool to insert this new stored procedure into the database.

•——CREATE PROCEDURE FROM——————*file_name*——————•

Figure 12-2 Syntax for the CREATE PROCEDURE FROM <file_name> statement

➲ **Example**

To create a procedure using multiple files:

```
dmSQL> CREATE PROCEDURE FROM 'proc1.ec';
dmSQL> CREATE PROCEDURE FROM '.\esql\sp\proc2.ec';
dmSQL> PROCEDURE FROM 'c:\users\jerry\sp\proc3.ec';
```

The previous examples show how to create stored procedures using dmSQL.

➲ **Alternatively, use JDBATool.**

1. Click the object **Stored Procedure** in the Tree.
2. Click the **Create** button. The **Introduction** window of the **Create Stored Procedure** wizard is displayed.
3. Import a stored procedure by selecting the **Import** button.
4. Selecting **Import** opens the **Open** window. Files can be imported from any source, including the SPDIR directory of other databases on the server or network drives. Select the desired file by typing in the path in the **File name** field, or browse through the directory tree until the correct path is found.

   Imported files must be ASCII format files that contain C++ code.
5. Select **Open** to open the file.
6. The **Create Stored Procedure** window will reappear if the imported file contains properly formatted (ASCII) text. Select **Save As** to store the stored procedure to another location, or select **OK** to compile and store the stored procedure in the database.

NOTE *If there are any errors when creating a stored procedure, they will be shown in the lower part of the window.*

NOTE

12.2 Executing Stored Procedures

You can invoke a stored procedure in dmSQL, a C program (ODBC or ESQL), another stored procedure, or using a trigger action.

dmSQL

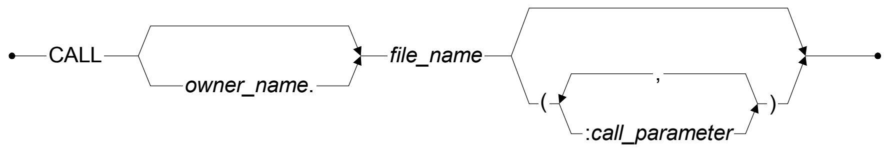

Figure 12-3 Syntax for the CALL statement within dmSQL

Example 1

To execute a stored procedure in dmSQL:

```
dmSQL> CALL p1 (3);                        // execute procedure p1
dmSQL> CALL SYSADM.p2 (5, ?);              // execute SYSADM's procedure p2
dmSQL/Val> 100;                            // input the value of parameter
dmSQL> ? = CALL SYSADM.p2 (5, 100);// execute procedure p2 and get
                                           // returned status
```

Example 2

If the procedure returns a result set, dmSQL automatically handles the output parameters and displays the result set on the screen. The result set appears on the screen as if you had typed a SELECT statement using dmSQL:

```
dmSQL> call a_phone('jeff');
       STATUS        PHONE
       =========================
       0       408-255-2689

dmSQL> call sel_all_phone;
       NAME              PHONE
       ============================
       Jerry            02-775-8615
       Jeff             408-255-2689
```

ESQL

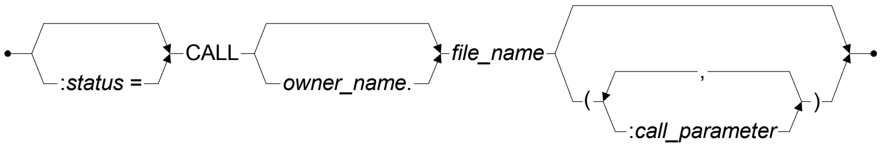

Figure 12-4 Syntax for the CALL statement within ESQL

➲ **Example**

To execute a stored procedure in ESQL:

```
EXEC SQL :s = CALL p1 (3);
EXEC SQL CALL SYSADM.p2 (5, :n2) INTO :nm;
EXEC SQL :s = CALL jack.p3 (:par1, 7) INTO :ret1, :ret2;
```

The syntax used in an ESQL program is similar to dmSQL. Use host variables to receive the status, output parameter, and result set values.

Executing Nested Stored Procedures

A stored procedure program is an ESQL/C program; a stored procedure inside another stored procedure may be invoked in exactly the same way it would be in an ESQL/C program. There is only one exception, regular ESQL programs cannot use the RETURNS keyword, but the stored procedure can use it when invoking another stored procedure.

Assume stored procedure **sel_all_phone** returns a multiple tuple result set. A regular ESQL program needs to use a cursor to fetch the tuples when invoking this procedure as shown in the last section. Another stored procedure **sp2** can use the same method to fetch tuples and examine the data. Also return the whole result set of the called stored procedure to the caller directly from within the current stored procedure.

➲ **Example**

To call a statement from within the sp2 stored procedure:

```
EXEC SQL RETURNS CALL sel_all_phone INTO :oName, :oPhone;
```


When a stored procedure returns another stored procedure's result set, the caller should have exactly the same result set list, or the first n result columns from the called procedure.

Executing Stored Procedures in ODBC programs

You can also call a stored procedure in an ODBC program by binding parameters for a procedure and using columns to return the result set. In an ODBC program, you can bind partial columns of the result set. After the procedure executes, output parameters are returned in the host variables. Use a fetch, like a SELECT command, to get the result set.

➲ **Example 1**

Procedure **proc1** declaration:

```
dmSQL> CREATE PROCEDURE proc1(CHAR(12) p1, CHAR(12) p2 OUTPUT) RETURNS INTEGER
r1;
{
EXEC SQL BEGIN CODE SECTION;
EXEC SQL SELECT c2 FROM t1 WHERE c1 = :p1 INTO :p2;
EXEC SQL RETURNS SELECT c1 FROM t2 INTO :r1;
EXEC SQL END CODE SECTION;
```

➲ **Example 2**

ODBC program that calls **proc1**:

```
SQLPrepare(cmdp,(UCHAR*)"call proc1(?, ?)", SQL_NTS);

strcpy(bpname, "12345");

SQLBindParameter(cmdp, 1, SQL_PARAM_INPUT_OUTPUT, SQL_C_CHAR, SQL_CHAR,
                      20, 0, &p1, 20, NULL);
SQLBindParameter(cmdp, 2, SQL_PARAM_INPUT_OUTPUT, SQL_C_CHAR, SQL_CHAR,
                      20, 0, &p2, 20, NULL);

SQLBindCol(cmdp, 1, SQL_C_LONG, &i, sizeof(long), NULL);
SQLExecute(cmdp);                                    /* get p2          */

while ((rc=SQLFetch(cmdp))!=SQL_NO_DATA_FOUND)       /* fetch result set */
```

Tracing Stored Procedure Execution

DBMaker provides trace functionality to help users trace the execution of stored procedures for debugging.

➲ **Example**

Using the TRACE command:

```
EXEC SQL TRACE ON;                          // Start TRACE
EXEC SQL SELECT c1 FROM t1 INTO :var1;
EXEC SQL TRACE ("var1 = %d\n", var1);     // TRACE the value of var1
EXEC SQL TRACE OFF;                         // END OF TRACE
```

Turn on and use the TRACE function to place variables for tracing and print messages. After the stored procedure executes, all trace information will be written to a file named **_spusr.log** in the **DB_SPLog** keyword directory found in the **dmconfig.ini** file on the client machine.

12.3 Dropping A Stored Procedure

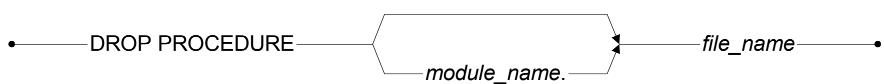

Figure 12-5 Syntax for the DROP PROCEDURE statement:

➲ **Example**

The first statement drops the stored procedure **proc1** the second statement drops stored procedure **user1.proc2**.

```
dmSQL> DROP PROCEDURE proc1;
dmSQL> DROP PROCEDURE user1.proc2;
```

12.4 Getting Procedure Information

➲ **Example 1**

To using dmSQL to get procedure information from the system table **SYSPROCINFO**:

```
dmSQL> SELECT * FROM SYSPROCINFO;
```

➲ **Example 2**

To use dmSQL to get procedure information from system table **SYSPROCPARAM**:

```
dmSQL> SELECT * FROM SYSPROCPARAM;
```

NOTE *ODBC functions* SQLProcedure() *and* SQLProcedureColumns() *are used to get procedure and parameter information for programs.*

12.5 Security

Only the owner or a user with DBA or higher authority can initially execute a stored procedure. Other users can execute the procedure when the execution privilege has been granted to them or a group that the user is a member of. Only owner or a user with DBA or higher authority can grant EXECUTE PROCEDURE privilege on a stored procedure for other users.

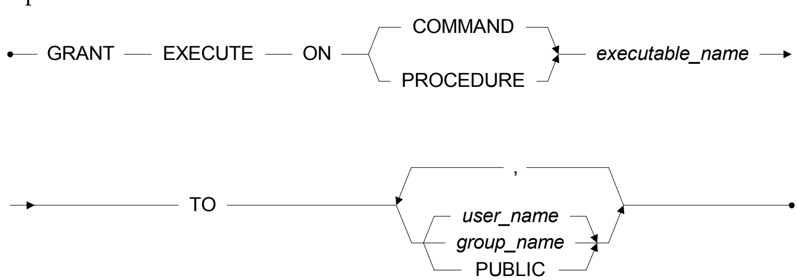

Figure 12-6 Syntax for the GRANT EXECUTE privileges statement

The owner or a user with DBA or higher authority can also revoke execute privilege on a stored procedure for other users.

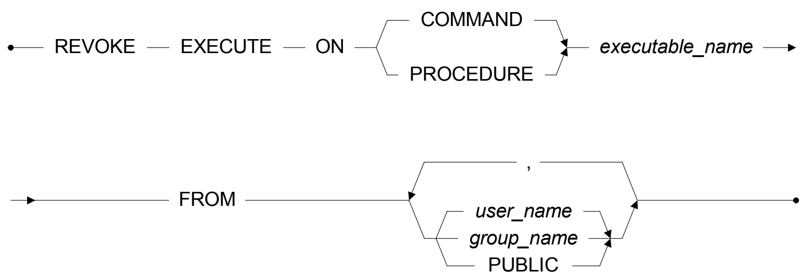

Figure 12-7 Syntax for the REVOKE EXECUTE privileges statement

➲ **Example 1**

user1 creates a stored procedure called **proc1** and grants the execute privilege to **user2** using dmSQL:

```
dmSQL> GRANT EXECUTE ON PROCEDURE proc1 TO user2;
```

➲ **Example 2**

user1 creates a stored procedure called **proc1** and grants the execute privilege to **PUBLIC** using dmSQL:

```
dmSQL> GRANT EXECUTE ON PROCEDURE proc1 TO PUBLIC;
```

➲ **Example 3**

user1 revokes the execute privilege from **user2** using dmSQL:

```
dmSQL> REVOKE EXECUTE ON PROCEDURE proc1 FROM user2;
```

➲ **Example 4**

user1 revokes the execute privilege from **PUBLIC** using dmSQL:

```
dmSQL> REVOKE EXECUTE ON PROCEDURE proc1 FROM PUBLIC;
```

13 Coding User-Defined Functions

DBMaker allows programmers to build their own user-defined functions (UDF). Once a UDF has been written in DBMaker, it is treated as a new built-in DBMaker function with the same usages. Creating a new user-defined function is straightforward and follows the general procedure outlined below:

➲ **To create a user defined function:**

1. Write a user defined function in C (UDF Interface)
   - **a)** Write the include statement
   - **b)** Write the function header
   - **c)** Write the arguments that the function passes.
   - **d)** Define allocated memory if necessary
   - **e)** Define an error code, if desired.
2. Build the dynamic link library for the UDF
3. Create the UDF in DBMaker, with the data array to be passed to the UDF.

13.1 UDF Interface

The first step in creating a UDF is coding it in C. The following sections give an example of a UDF in C, and describe each of the elements of the code that are particular to a DBMaker UDF.

Example

If you want to create a new UDF, **INT2STR()**, to convert integer data to a string, you should build a dynamic link library to include it.

```
dmSQL> SELECT INT2STR(c1) FROM t1;      // c1 is integer type
```

The following C source code, *template.c*, gives a snapshot of code of the **INT2STR()** UDF:

```
#include <memory.h>
#include <string.h>
#include <stdio.h>
#include "libudf.h"

/* Transfer integer type to string type */
#ifdef WIN32
__declspec(dllexport)
#endif
int INT2STR(int narg, VAL args[])
{
  char *ptag;
  int len;
  char p1[11];
  int rc;

  if (args[0].type != NULL_TYP)
    {
    sprintf(p1, "%d", args[0].u.ival);
    len = strlen(p1);
    if (rc = _UDFAllocMem(args, &ptag, len))
        return rc;
    memcpy(ptag, p1, len);
    args[0].type = CHAR_TYP;
    args[0].len = len;
    args[0].u.xval = ptag;
    }
  return _RetVal(args, args[0]);
}
```

Including libudf.h

DBMaker defines some constants, data types and common interfaces, which are needed in UDF coding.

Programmers should include **libudf.h** before any UDF coding:

```
#include "libudf.h"
```

Passing Parameters

The arguments of a UDF used in an SQL command are packaged into the **args** parameter of the UDF coded in C language. Through the **args** array, a UDF gets the input data. **args** is also called the *UDF control block*, which is always used as the first argument of the common interface provided by DBMaker. Some common interfaces, such as the BLOB Common Interface, will be introduced later.

Each UDF header in a C function should follow the form:

```
int FUNCTION_NAME(int narg, VAL args[])
{
...
}
```

NOTE *args[] points to an array. Functions passing only one argument should use the pointer form: *args.*

narg specifies how many arguments the function passes. For example, if a UDF **MYSUBSTRING** (**c1**, **c2**, **c3**) is called in an SQL command, **c1** information is passed by *args[0]*, **c2** by *args[1]* and **c3** by *args[2]*. The value of *narg*, specifying the array size, is 3.

➲ **Example 1**

Using the value of **c1** as 'abcdefghijklmn', *args[0]* would be:

```
args[0].type   = CHAR_TYP
args[0].len    = 14
args[0].u.xval = "abcdefghijklmn"
```

➲ **Example 2**

Using the value of **c2** as integer 30, *args[1]* would be:

```
args[1].type   = INT_TYP
args[2].len    = 4
args[3].u.ival = 30
```

In addition to CHAR_TYP and INT_TYP, BIN_TYP, FLT_TYP, OID_TYP, BLOB_TYP, DEC_TYP and NULL_TYP constants are defined in **libudf.h**:

```
#define BIN_TYP   0x0000                   /* bit string     data type*/
#define CHAR_TYP  0x1000                   /* character      data type*/
#define INT_TYP   0x2000                   /* integer        data type*/
#define FLT_TYP   0x3000                   /* floating point data type*/
#define OID_TYP   0x4000                   /* OID            data type*/
#define BLOB_TYP  0x5000                   /* BLOB           data type*/
#define DEC_TYP   0x6000                   /* decimal        data type*/
#define NULL_TYP  0xF000                   /* set if column is null   */
```

➲ **Example 3**

Through NULL_TYP, the programmer can know whether the input data is NULL:

```
if (args[0].type == NULL_TYP)
{
    /* input data is NULL */
}
else
{
   /* input data is not NULL */
}
```

The complete data structure of VAL, as defined in **libudf.h**:

```
typedef struct tVAL {
   i16 type;                   /* data type          */
   i15 len;                    /* data length        */
   union {
      i31    ival;             /* long integer data  */
      i15    sival;            /* short integer data */
      double fval;             /* double data        */
      float  sfval;            /* float data         */
      dec_t  dval;             /* decimal data       */
      char   *xval;            /* pointer to data    */
   } u;
} VAL;
```

The structure *dec_t*, used for DECIMAL type, in **libudf.h**:

```
typedef struct
{
    i8 pre;
    i8 sca;
    i8 dgt[9];
    i8 exp;
} dec_t10;
typedef dec_t10 dec_t;
```

A UDF not only passes input data through *VAL* type, but also returns output data through it. How to return data is discussed later.

Allocating Memory Space

In C functions, you may need to allocate memory and free it before leaving the function. Returned values, such as a string or temporary BLOB ID, need to allocate memory, hold it in the UDF function, and have DBMaker assist in freeing memory space.

➲ **Example**

In the following example of a UDF **UDFAllocMem**, *arg* is the UDF control block, *ppt* is the pointer to get the allocated memory block, and *nb* is the desired allocated size. This function allocates memory and holds it until DBMaker takes care of it:

```
int _UDFAllocMem(VAL *arg, char **ppt, int nb);
```

DBMaker knows to free the memory after a result is returned by using *args[0].u.xval*, a pointer to memory space allocated by **_UDFAllocMem**().

```
if (rc = _UDFAllocMem(args, &ptag, 10))
   return rc; /* return error code */
memcpy(ptag, "0123456789", 10);
args[0].type = CHAR_TYP;
args[0].len = len;
args[0].u.xval = ptag;
```

Returning Results

There are two types of returned values: one is an error code and the other is the result of the UDF through the argument type VAL. Error codes are returned to DBMaker but their values are hidden from the user; only an error message will be displayed. The following describes how error codes are returned.

The header of UDF in a C function follows the form:

```
int FUNCTION_NAME(int narg, VAL args[]);
```

If **FUNCTION_NAME**() returns a non-zero value there is something wrong, if a 0 is returned it means that the function worked properly.

Before returning from the UDF, call **_RetVal**() to pass the imported result from the UDF to DBMaker with the following declaration:

```
int _RetVal(VAL *arg, VAL rtn);
```

The first argument *arg* is the UDF control block, and the second one *rtn* is the value returned. The following code returns integer 30:

```
int rc;                      /* error code */
VAL rtn;
rtn.type   = INT_TYP;
rtn.len    = 4;
rtn.u.ival = 30;
rc = _RetVal(arg, rtn);      /* pass result back to DBMaker         */
return rc;                   /* return error code (0 means no error) */
```

13.2 Building UDF Dynamic-Link Library

DBMaker provides a library **dmudf.lib** to link with the UDF source file to build the dynamic-link library. Since the dynamic-link library is different on Microsoft Windows and UNIX environments, both cases are discussed separately.

DLL in Microsoft Windows Environment

DBMaker also provides the **template.c** source code in the **/udf_templates** directory and the template make files **udf42.mak** (for Microsoft VC++ 4.2 version), **udf50.mak** (for Microsoft VC++ 5.0 version), or **udf60.mak** (for Microsoft VC++ 6.0 version) for WIN32 users to reference. Users can follow the format of a template C source file to write their UDF.

➲ **In the following statements, the udf60.mak is used.**

1. Ensure where to include the **dmudf.lib** file and then use the IDE that Visual C++ provides to modify the required changes.
2. Copy **udf60.mak** template make file into the desired directory and rename it with a make file name.
3. Choose <File> -> <Open Workspace> to open the make file project workspace.
4. In the <Project Workspace>, choose <File View>, click **template.c** to remove it, press Delete
5. Choose <Project> item in the tool bar, choose < Add to project >, <Files>, and, insert your own .c file into the make file of the project workspace.
6. In the <Project> -> <Settings>, choose WIN32 Debug for this example. In the Project Settings <General>, you can change output directories. In the template make file, set 60Deb as the intermediate and output directories.
7. In the Project Settings <Link> item, in Category item <General>, change the output .dll file name directly in <Output file name>. Also, change the link path of the *dmudf.lib* file DBMaker supports in the <Object/Library modules> to the working directory.

After finishing the above, you can build your own **dll make** file. Using similar steps, you can also build a WIN32 Release version dll file.

Users of VC++, can also create a dll make file using the same steps and setting the structure member alignment to be 4 bytes. In the VC ++ 6.0 IDE project workspace, choose the C/C++ menu item, then in the **Category** dialog box, choose <Code Generation>. You can find the structure member alignment option, and then choose 4 bytes as the result.

Use the make file template to note the setting when writing a **collect dll**. If you do not want to use **template.c** as the default C filename within the make file, remove ***template.c*** from ***udf60.mak*** and insert your C file into the ***udf60.mak*** project workspace.

➲ **Example**

In the DBMaker **template.c**, remember to include the **libudf.h** file provided, and to export your functions. Use the export function method from the VC++ programmer guidebook or the following:

```
__declspec(dllexport) datatype YOUR_FUNCTION_NAME( ...... )
```

Alternatively, create a **def file** in the project workspace to export your functions and note that the function name for the UDF must be in UPPER CASE, in C source code.

After finishing the above, you can build a **debug/release version dll** file, thus creating a ***udf60.dll*** file.

UDF so File in UNIX

A **so** file, or UNIX dynamic library, can be created.

➲ **Example**

Write UDF C source code, in the example the file is named **udf.c**. After completion, use the UDF function in a UNIX based OS

```
$ cc –c udf.c
$ ld –o libudf.so udf.o –lm
$ dmsqlt
dmSQL> CREATE FUNCTION libudf.INT2STR(INT) RETURNS CHAR(10);
```

NOTE *The options of the* ***ld*** *command in the above example can vary in UNIX. It may be –G, –shared, or something else. Please refer to your UNIX manual or man pages to check how to use the* ***ld*** *command in building a shared library.*

13.3 Creating, Using, and Dropping UDF

The next step for a user-defined function is to create it within DBMaker. The following sections outline the syntax for creating, querying, and dropping a UDF.

Creating a UDF

➲ **Syntax**

```
dmSQL> CREATE FUNCTION <udf_dll_name.function_name> (<function_datatype>) RETURNS
<function_output_datatype>;
```

Querying a UDF

➲ **Syntax**

```
dmSQL> SELECT <function_name>(<related_table_column_name>)
FROM <related_table>;
```

Dropping a UDF

➲ **Syntax**

```
dmSQL> DROP FUNCTION <function_name>;
```

Example

The following demonstrates how to use a UDF file.

➲ **Example 1**

Using a database named **DMDEMO** containing a table, **t1,** with the table schema, **c1 INT, c2 CHAR(10)**:

```
dmSQL> SELECT * FROM t1;
c1          c2
==================
10         1
20         2
30         3
```

```
3 rows selected
```

Using the example **template.c** DBMaker supports, we can now build a **udf60.dll** successfully.

In the **dmconfig.ini** file, add one line to the **DMDEMO** section:

```
[DMDEMO]
DB_DBDir = D:\UDFDEMO
DB_FODIR = D:\UDFDEMO\FO
DB_LBDIR = D:\UDF\60Deb   ; add this line
```

For more information on **DB_LbDir**, refer to Appendix A *Keywords in dmconfig.ini*. Set **DB_LbDir** or place the **udf60.dll** in the **<DBMaker home directory>\shared\udf**, since it is the UDF default directory.

➲ Example 2

Start the database **DMDEMO**, and then create the UDF function. In the example, the <*udf_dll_name*> is **udf60**,the <*function_name*> is **INT2STR**, <*function_datatype*> is INT, and <*function_output_datatype*> is CHAR(10):

```
dmSQL> CREATE FUNCTION udf60.INT2STR(INT) RETURNS CHAR(10);
```

The UDF function **INT2STR** returns the following results. The <*function_name*> is **INT2STR**, <*related_table_column_name*> is **c1** according to the schema of **t1** and the <*related_table*> is **t1**:

```
dmSQL> SELECT INT2STR(c1) FROM t1;

INT2STR(c1)
===========
10
20
30

3 rows selected
```

➲ Example 3

Another UDF function, e.g. **STR2INT()**, in the same dynamic-link file:

```
dmSQL> CREATE FUNCTION udf60.STR2INT(CHAR(10)) RETURNS INT;
```

```
dmSQL> SELECT STR2INT(c2) FROM t1;

 STR2INT(c2)
===========
1
2
3

3 rows selected
```

➲ **Example 4**

When dropping a UDF function, simply drop the UDF function name, there is no need to attach the **UDF dll** name. When dropping a UDF function, wait until the database has terminated, then the UDF function will be cleaned up. Before the database is terminated, the function will continue to exist.

```
dmSQL> DROP FUNCTION INT2STR;
```

13.4 UDF BLOB Common Interface

Today, multimedia is important and useful to users. DBMaker supports a common interface to access BLOBs using a file handle method, so programmers can easily write UDFs for BLOB type data. FILE, LONG VARCHAR, and LONG VARBINARY are the data types used to store BLOB data in a database.

Many of the new features in DBMaker need a temporary BLOB to process temporary results. DBMaker supports temporary BLOBs for programmers to write a UDF more easily. A programmer can open a permanent BLOB, read the data, execute a conversion function or something else, save the result in a new temporary BLOB and return it back in a UDF. The API fetches this temporary BLOB as a normal BLOB column.

BLOB Common Interface Functions

DBMaker provides BLOB common interface functions for programmers to write UDFs. A DBA should set the **DB_FODir** in the **dmconfig.ini** file for the temporary BLOB file before starting a database. A temporary BLOB will be created in an external file in the **DB_FODir** directory, with the file name format "**__??????.TMP**", where "?" represents one character of either [0-9, A-Z]. All file names matching the format will be deleted when the database is shut down and restarted.

_UDFBbOpen()

Opens a BLOB using **bbObj** and returns a handle through **pHandle**. **bbObj** can be retrieved by **Arg**[**i**], using the **BLOB** with the input argument of the UDF. The function returns 0 if it successfully opens the **BLOB**, otherwise an error code will be returned:

```
int _UDFBbOpen(VAL *Arg, BBObj bbObj, i31 *pHandle);
```

_UDFBbRead()

Reads the BLOB that belongs to the handle. Before calling this function, allocate a buffer, (**pBuf**), with **szBuf** using the function **_UDFAllocMem**() to get the read data. The returned data will be stored in **pBuf** and the size actually read is in **szRead**. If **szBuf** is non-positive, no characters are read:

```
 int _UDFBbRead(VAL *Arg, i31 handle, i31 szBuf, i31 *szRead, char
*pBuf);
```

_UDFBbSeek()

This function is used to set the position of the next output operation in a **BLOB**. The new position is at the offset bytes from the beginning, the current position, or the end of the file, according to the **ptrname** using the **SEEK_BB_BEG**, **SEEK_BB_CUR**, or **SEEK_BB_END** value defined in **libudf.h**. The function only works between the period of **_UDFBbOpen**() and **_UDFBbClose**(), but not **_UDFBbCreate**() and **_UDFBbClose**():

```
int _UDFBbSeek(VAL *Arg, i31 handle, i31 offset, i16 ptrname);
```

_UDFBbCurOffset

The function returns the current position of an open **BLOB** or the offset in a **BLOB**:

```
int _UDFBbCurOffset(VAL *Arg, i31 handle, i31 *pOffset);
```

_UDFBbClose()

Closes the **BLOB** opened by **_UDFBbOpen**() or created by **_UDFBbCreate**():

```
int _UDFBbClose(VAL *Arg, i31 handle);
```

_UDFBbCreate()

Creates a temporary **BLOB** and returns a handle for **_UDFBbWrite**(). The caller should prepare the space for the **BBObj** structure pointed to by **pBbObj** and written by **_UDFBbCreate**(), **_UDFBbWrite**() and **_UDFBbClose**(). **BBObj** is used to identify this temporary **BLOB**. For example if you want to delete the temporary **BLOB** called **_UDFBbDrop**() using the **BBObj** argument.

If successful, **pHandle** will return a **BLOB** handle similar to the handle of the opened file written by **_UDFBbWrite**() and closed by **_UDFBbClose**().

Alternatively, specify the temporary **BLOB** to be created in file (**BB_TEMP_FO**)or in memory (**BB_TEMP_MEM**). If the caller specifies the temporary **BLOB** in memory, it does not mean that the temporary **BLOB** will be created in memory - a memory limitation may prevent this. The temporary **BLOBs** in memory might be converted to files by the operating system if the original temporary **BLOBs** in memory or the input data are over the size limit. Programmers should not depend on this feature when coding.

The function returns **0** if it is successful and an error code will be returned otherwise.

Before reading the new temporary **BLOB**, you must close it using **_UDFBbClose**(), then reopen it using **_UDFBbOpen**(). **_UDFBbSeek**() cannot be used on temporary **BLOBs** unless they are closed and reopened for reading :

```
int _UDFBbCreate(VAL *Arg, BBObj *pBbObj, i31 *pHandle, i31 Opt);
```

_UDFBbWrite()

After using **_UDFBbCreate**() to make a temporary **BLOB**, write data to it using **_UDFBbWrite**(). The handle is from **_UDFBbCreate**(), **pBuf** points to input data and its length is **szBuf**. The function returns **0** if it is successful, otherwise, an error code will be returned:

```
int _UDFBbWrite(VAL *Arg, i31 handle, i31 szBuf, char *pBuf);
```

_UDFBbDrop()

Normally you do not drop a temporary **BLOB** if it will be returned from a **UDF**; the system will control its life cycle. If you do not return the created **BLOB**, you'd better use this function to drop the temporary **BLOB**. This function cannot work on a permanent **BLOB**; doing so will return the **ERR_BLOB_INV_BLOB** error. The function returns **0** if it is successful, otherwise, an error code will be returned:

```
int _UDFBbDrop(VAL *Arg, BBObj bbObj);
```

_UDFBbSize()

This function returns the data size of a **BLOB** by **pLen**. **BbObj** can be a permanent BLOB or a temporary BLOB. The function returns 0 if it is successful, otherwise, an error code will be returned:

```
int _UDFBbSize(VAL *Arg, BBObj bbObj, i31 *pLen);
```

Example

The following demonstrates how to create the user-defined function, **MYCONVERT** with input in varchar format and output as a temporary BLOB.

➲ **To create the user-defined function, MYCONVERT:**

1. Build a dynamic library in UNIX using **myudf.c**, (the source code follows later):

```
cc -g -c myudf.c
ld -G -o myudf.so myudf.o
```

2. Start the database.

3. At the dmSQL prompt, enter:

```
dmSQL> CREATE FUNCTION myudf.myconvert(VARCHAR(100))  // input string
    2>  RETURNS LONG VARCHAR;                         // output BLOB
dmSQL> SELECT myconvert(c1) FROM mytable;             // output temp BLOB
```

The source code for the UDF **MYCONVERT**:

```
#include "libudf.h"
int MYCONVERT(int nArg, VAL args[])
{
  int     rc = 0, trc;              /* return code                    */
  BBObj   tmpobj;          /* output temp BLOB               */
```

```
  i31     handle;          /* handle of created temp BLOB     */
  boolean fgCreate = false;        /* temp BLOB has been created?     */
  char   *pInData, pOutData[4096];/* input/output data buffer        */
  i31     nInData, nOutData;       /* input/output data buffer length */

  if (args[0].type == NULL_TYP)
    goto cleanup;

  pInData = args[0].u.xval;        /* get input data                  */
  nInData = args[0].len;           /* input data length               */

  /* create a temp BLOB in file */
  if (rc = _UDFBbCreate(args, &tmpobj, &handle, BB_TEMP_FO))
    goto cleanup;
  fgCreate = true;

  /* any real processing function */
  RealConvert(pInData, nInData, pOutData, &nOutData);

  /* write result data to temp BLOB */
  if (rc = _UDFBbWrite(args, handle, nOutData, pOutData))
    goto cleanup;

  /* close created temp BLOB ( temp BLOB is still alive) */
  if (rc = _UDFBbClose(args, handle))
    goto cleanup;

  args[0].type = BLOB_TYP;
  args[0].len = sizeof(BBID);
  args[0].u.xval = (char *)&tmpobj;
                /* _RetVal() does a copy from this local buffer  */

cleanup:
  if (rc)
    {
    /* error handle */
    if (fgCreate)
      {
      _UDFBbClose(args, handle);      /* close created temp BLOB     */
```

```
      trc = _UDFBbDrop(args, tmpobj); /* drop it because of failure */
      if (trc > rc)
        rc = trc;
      }
    return rc;
    }
  else
    return _RetVal(args, args[0]);

}/* MYCONVERT() */
```

Troubleshooting Errors

Use the following to troubleshoot errors when writing a BLOB UDF using the BLOB common interface.

ERROR (327): THE BLOB COLUMN IS NOT OPENED OR CREATED YET

The function must use **_UDFBbOpen**() to open the BLOB or **_UDFBbCreate**() to create a new temporary BLOB, before using other BLOB function interfaces.

ERROR (328): THE OFFSET OF BLOB COLUMN IS INVALID

When a UDF using **_UDFBbSeek**() seeks to offset by a length greater than the length of the BLOB.

ERROR (331): THIS BLOB WAS NOT IN CREATED STATE

_UDFBbWrite() can only work on a temporary BLOB created by **_UDFBbCreate**() and must not be closed. For example, if you use it on BLOB opened by **_UDFBbOpen**(), this error will occur.

ERROR (330): THIS BLOB WAS NOT IN OPENED STATE

_UDFBbRead() can only work on a BLOB (including temporary BLOBs) opened by **_UDFBbOpen**().

ERROR (332): THE BLOB OBJECT IS NOT CLOSE YET

Whenever **_UDFBbOpen**() or **_UDFBbCreate**() are used to open a BLOB, programmers should call **_UDFBbClose**(), to close the opened BLOB.

ERROR (322): NO FILE OBJECT DIRECTORY IN CONFIGURATION FILE; CANNOT INSERT FILE OBJECT

If temporary BLOBs are used, the keyword **DB_FODir** in the ***dmconfig.ini*** file must be set. If not set, attempting to create a temporary BLOB may fail and this error occurs.

13.5 UDF related dmconfig.ini keywords

DB_StrSz

In addition to **DB_LbDir** and **DB_FODir**, there is also a related keyword **DB_StrSz** in **dmconfig.ini** file

```
DB_STRSZ=<value>
```

This keyword indicates the length of returned data of the STRING type, used only by user-defined function (UDF). Since UDFs can only return data of a fixed size, this keywords can limit the size of STRING data, in order to avoid receiving strings that are to long. The default value is 255, and the valid range is from 1 to 4096. It can be used on a client or server, the client has a higher priority.

14 Database Recovery, Backup, and Restoration

In every database management system, the possibility of a hardware or software failure always exists. A DBMS may fall victim to failures without warning. After a failure occurs, a DBMS should have some method of recovering the information. This is one of the main advantages a DBMS has over the old file-based systems they replaced.

DBMaker incorporates advanced data protection features to prevent data loss and downtime due to failures. These features allow DBMaker to ensure the reliability of a database and the consistency of data by providing recovery, backup, and restoration features.

14.1 Types of Database Failures

Database failures can be divided into two types: system failures and media failures. When either of these types of failure occurs, there is the possibility of data inconsistency or data loss in a database. A DBMS should provide facilities for recovering from failures and for replacing a damaged database with a backup copy.

System Failures

A system failure, known as an instance failure, is a failure from the main memory in a computer system. System failures may be caused by a power failure, an application or operating system crash, a memory error, or other reason. The result is the unexpected termination of DBMS.

Applications and active transactions can terminate abnormally when a system failure occurs. Since the exact state of a transaction in progress or a transaction that has not been completely written to disk cannot be reliably be determined after a system failure, these types of transactions require recovery. The most common method of protection against system failures is the use of a transaction log, or a journal file.

Media Failures

Media failure (also known as disk failure) is a failure of the disk storage system of a computer system. Media failures are usually caused by physical trauma to the disk itself, such as a head crash, fire, or exposure to vibration or g-forces outside its physical operating limits.

There is nothing to prevent the loss of data on an affected disk when a media failure occurs. One or more files may be physically damaged because of the failure, requiring restoration of the database. However the database can be successfully restored if the database provides backup and restoration facilities.

14.2 Recovery from Database Failures

The goals of recovery after a database failure are to ensure committed transactions are reflected in the database, ensure uncommitted transactions are not reflected in the database, and to return to normal operation as quickly as possible while insulating users from problems caused by the failure.

DBMaker uses journal files and checkpoints to achieve these goals. The journal files and checkpoints work together to ensure that all transactions are recovered in as short a time as possible, with as little effect on users as possible.

Journal Files

Journal files provide a real-time, historical record of all changes made to a database, and the status of each change. In the event of a system failure, the historical record of changes maintained in the journal file allows DBMaker to recover and redo changes made by transactions that completed but were not written to disk, or undo changes made by transactions that terminated abnormally.

If a database is running in backup mode, the journal files will also store additional information that DBMaker can use to restoration. This information will remain in the journal files until they are backed up; after you back up the journal files DBMaker will free this space for use by new transactions.

During the restoration process, DBMaker will add the information from the backup journal files to a backup copy of the database. Therefore, only the journal files that contain the changes made to the database between full backups need to be backed up.

Checkpoint Events

A checkpoint is a system event in which the database is brought to a clean state. DBMaker writes all journal records and all dirty data pages from its internal memory buffers to disk, and reclaims journal blocks that are no longer required for backup or recovery purposes. DBMaker can reclaim journal blocks that contain non-active transactions that completed before the start of the oldest active transaction.

Startup time after an instance failure is reduced after taking a checkpoint. DBMaker writes the time of the last checkpoint and a list of all transactions active at the time of the checkpoint to the journal file header. During database recovery, DBMaker uses this information to determine which transactions should be undone, which should be redone, and which should be ignored.

DBMaker will automatically take a checkpoint when the journal files are full to try to reclaim some journal blocks to reuse. If the checkpoint cannot reclaim enough space to complete the current transaction, the transaction will be aborted. DBMaker will also automatically take a checkpoint when the database starts and shuts down, and when an online backup is performed.

Database administrators can initiate a checkpoint manually by executing the CHECKPOINT command. The optimal interval between two checkpoints depends on the frequency of activity in the database, the average size of transactions, and the size and number of journal files. Since these factors may vary significantly from database to database, the optimal interval is best determined through experience. Manual checkpoints reduce the amount of time required to start, terminate, and backup a database, as well as the possibility that a full journal will be encountered.

Checkpoints may require a significant amount of time to complete, depending on the size and number of transactions since the last checkpoint. Any transactions that are active when a checkpoint occurs need to wait for DBMaker to calculate which journal records it can reclaim, but do not need to wait while DBMaker actually writes journal records and dirty data pages to disk.

Recovery Steps

DBMaker provides support for automatic recovery when the database is started after a system failure or when an error occurs during startup. During the recovery process, DBMaker always performs two separate steps: redo and undo.

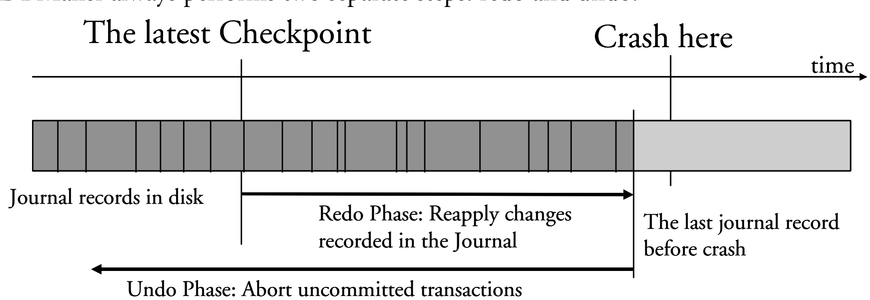

The first step in the recovery process is to redo (or reapply) all changes made to the database that are recorded in the journal. This step is necessary since it is possible for a transaction to have completed before the system failure, without having all the changes made by the transaction written to the database. However, these changes *are* stored in the journal, and can be written to the database during this step. At the end of this step, the database contains the changes made by all committed transactions, as well as the changes made by all uncommitted transactions.

The second step in the recovery process is to undo (or rollback) all the changes made by transactions that were not completed before the system failure occurred. This step is necessary since the exact state of a transaction in progress cannot be reliably determined in the event of a system failure. These incomplete transactions must be removed since a transaction is self-contained by definition and must either complete successfully and change the data, or fail and leave the data unchanged. At the end of this step, the database contains the changes made by all committed transactions, but does not contain any changes made by uncommitted transactions.

DBMaker also provides support for starting a database after a media failure or after a system failure, which causes inconsistencies in a database that cannot be repaired during the automatic recovery process. In these cases, the database will fail to start and you would normally need to restore your database from a backup copy. However, if you have never backed up your database, you can force the database to start by setting the forced-start mode using the **DB_ForcS** keyword in the **dmconfig.ini** file. This will allow you to start the database and unload the unaffected data. For more information on the forced-start mode, see “Forcing Database Startup”.

Forcing Database Startup

DBMaker automatically performs recovery operations if errors occur when a database starts normally. If the database cannot start up, there may be some disk errors. Disk errors require the database be restored from the most recent backup to repair it. If the database has no backups and cannot start, use the *forced startup* mode provided by DBMaker.

DBMaker supplies a forced startup option for this type of situation. To set the forced startup mode on, use the **DB_ForcS** keyword in the **dmconfig.ini** file. Setting this keyword to 1 enables forced startup mode, and setting it to 0 disables it. When forced startup mode is on, DBMaker will skip errors when starting the database.

If the database still cannot be started, there is one remaining alternative provided in the procedure below. However, before performing this procedure, backup all data and journal files.

➲ **To start a database when it will not start in force start database mode:**

1. Set the Forced Startup Mode to off in **dmconfig.ini** (**DB_ForcS** = 0).
2. Set the Start Mode to New Journal Mode in **dmconfig.ini** (**DB_SMode** = 2).
3. Restart the database.
4. Reset Start Mode back to normal in **dmconfig.ini** (**DB_SMode** = 1)

DBMaker provides the option to use a new journal to force the database to start without any recovery operations. Therefore, if errors serious enough to prevent the database from starting have occurred, the database may be in an inconsistent state.

After starting the database with this method, check the consistency of the database. For more information on database consistency checking, refer to section 6.12 *Checking Database Consistency*.

14.3 Types of Backups

Backups are used to protect a database from media failures or other media errors. After a media failure, one or more database files may be physically damaged and unusable. Use the most recent backup to replace the damaged files and reconstruct a database.

Database backups consist of backup sequences. One backup sequence consists of one full backup and any incremental backups that were performed after the full backup.

Full Backups

A full backup is any backup that creates a copy of all data and journal files, providing a copy of the entire database at one point in time. A backup copy of the **dmconfig.ini** file can be created as well, preserving any custom configuration settings there may be for a database. The database administrator may perform a full backup while the database is online or offline.

Full backups archive the entire database, therefore requiring a large amount of storage space. However, a database can be restored relatively quickly using a full backup. A full backup can be used to restore a database to the point in time the full backup was performed.

A valid full backup will be assigned a full backup ID. The full backup ID is a time / date stamp. The backup ID is used to ensure that full backups and incremental backups are associated in a backup sequence; all incremental backups between the current valid full backup and the next belong only to the current valid full backup. Trying to restore incremental backups against previous (and any other) sequences will fail. Backup sequences are managed by DBMaker. Repairing a database, restoring a database, starting a database in new journal mode, or changing the backup mode will require a new valid full backup.

There are three primary methods of performing full backups. The first is by using the backup server, and is discussed in more detail in section 14.6 *Backup Server*. Full backups by backup server may be performed with dmSQL or with the JServer Manager utility. The second method is full backup interactively. Full backup interactively does not require that the Backup Server be started. JServer Manager is the recommended method of performing this type of full backup. For directions on how to perform a full backup interactively, refer to the *JServer Manager User's Guide*. The third method for performing a full backup is offline full backup. Refer to section 14.5 *Offline Full Backups* for more information on how to perform an offline full backup.

Incremental Backups

An incremental backup is any backup that creates a copy of only the journal files that have changed since the last incremental or full backup. Incremental backups may only be performed after a full backup has been performed. Performing a new full backup starts a backup sequence. Subsequent incremental backups are part of that sequence and may not be used with any other full backup. The incremental backup file sequence provides a copy of the changes made to the database since the last full backup. The database administrator can perform an incremental backup only while a database is online.

Incremental backups archive only journal files, so they require only a small amount of storage space. However, it may take more time than a full backup to restore a database since the DBMS must take the time to rollover all transactions in the backup journal files. Use the incremental backups together with a full backup to restore a database to any point in time between the previous full backup and the time the last incremental backup was completed.

There are two methods for performing Incremental Backups (there is a third method Incremental backup to current, that is considered a different type of backup). The first method is Incremental backup by backup server. Backup server must be started on the database to be able to use this method. For more information on performing incremental backup by backup server, refer to section 14.6 *Incremental Backups*. The second method is incremental backup interactively. This type of Incremental backup does not require that backup server be started. The recommended method of performing incremental backup interactively is with the JServer Manager utility. Incremental backup interactively is explained in the *JServer Manager User’s Guide*.

Offline Backups

An offline backup is any backup that must be performed after a database has been shut down. The database administrator must schedule a time to shut down the database, and notify all users so they can disconnect from the database. Offline backups can be inconvenient for users, since they must remember to complete all active transactions and disconnect from the database before it is shut down. The database administrator can perform only full backups while offline.

A DBMS does not need to provide the capability to backup a database offline, since you can backup the database with operating system commands after it is shut down. The database administrator may perform an offline backup using this method, or by using JServer Manager, an easy-to-use graphical tool that performs offline backups without resorting to using operating system commands.

Online Backups

An online backup is any backup that is performed while a database is running. The database administrator does not have to shut down the database, and users do not need to disconnect. Online backups are more convenient for users, since no action is required on their part. The database administrator can perform full and incremental backups while online.

A DBMS must provide the capability to backup a database online, since it is still running and still has users connected. DBMaker allows for online backups to be performed manually using dmSQL and operating system commands, but also provides JServer Manager, an easy-to-use graphical tool that allows online backups to be performed without resorting to operating system commands.

Online Incremental to Current Backups

DBMaker also supports an additional backup type known as online incremental to current.

The difference between an online incremental backup and an online incremental to current backup in a database with multiple journal files is minor, but important. In an online incremental backup DBMaker will backup all journal files that have been used since the last backup, excluding the active journal file. In an online incremental to current backup DBMaker will backup all journal files that have been used, including the active journal file. This means that an online incremental backup can restore a database up to the point in time the last committed transaction was written to the last full journal file, while an online incremental to current backup can restore a database up to the point in time the active journal file was backed up.

In a database with only a single journal file, an online incremental backup and an online incremental to current backup are the same. In this case, the only journal file is the active journal file. DBMaker will backup this single journal file in both types of incremental backup.

Online Incremental to current backups may be performed with the JServer Manager Utility. For directions on how to perform an incremental backup to current, refer to the *JServer Manager User's Guide*.

14.4 Backup Modes

Backup mode determines whether DBMaker can perform online incremental backups, and the type of data that will be backed up during an incremental backup. It also determines when DBMaker will free journal blocks that belong to inactive transactions for use by other transactions. DBMaker has three backup modes: NONBACKUP, BACKUP-DATA, and BACKUP-DATA-AND-BLOB.

Backup Mode	Tablespace Backup Mode	User Defined Tablespace (Data)	User Defined Tablespace (BLOB)	System Tablespace (Data and BLOB)
No Backup		No	No	No
Backup Data		Yes	No	Yes
Backup Data and BLOB	Backup BLOB Off	Yes	No	Yes
	Backup BLOB On	Yes	Yes	Yes

NONBACKUP Mode

NONBACKUP mode provides no protection for any data that was inserted or updated since the last full backup. In this mode, online incremental backups cannot be performed. A database can use the journal to fully recover from instance failure, but a media failure may result in loss of data. Journal blocks not in use by an active transaction can be reused immediately after a checkpoint, but once they are overwritten, the database may only be restored to the point in time of the last full backup.

BACKUP-DATA Mode

BACKUP-DATA mode provides protection for data (excluding BLOB data) that was inserted or updated since the last full backup. In this mode, a database administrator can perform an online incremental backup, but only non-BLOB data will be stored in the backup files. A database can use the journal to fully recover from instance failure, and can partially recover from media failure. Although the last backup can be used to restore the database to the point in time of the media failure, any changes to BLOB data will be lost. Journal blocks not in use by an active transaction can only be reused after a checkpoint has taken place *and* the journal file has been backed up.

BACKUP-DATA-AND-BLOB Mode

BACKUP-DATA-AND-BLOB mode provides protection for all data (including BLOB data) that was inserted or updated since the last full backup. In this mode, a database administrator can perform an online incremental backup, and all data will be stored in the backup files. A database can use the journal to fully recover from instance failure, and can fully recover from disk failure. The last backup may be used to completely restore the database to the point in time of the media failure, including all BLOB data. Journal blocks not in use by an active transaction can only be reused after a checkpoint has taken place and the journal file has been backed up.

Tablespace BLOB Backup Mode

DBMaker normally applies the backup mode setting to the entire database; this means all tablespaces in the database will be in the same backup mode. If the database is in BACKUP-DATA-AND-BLOB mode, DBMaker will record all changes to data (including BLOB data) in the journal. Recording BLOB data in the journal can quickly exhaust journal space, producing frequent backups and large backup file sizes.

This may be necessary if all BLOB data is critical, but in many cases, non-critical BLOB data may be backed up at the same time. Situations like this make it difficult for the database administrator to decide which backup mode you should choose. To prevent this type of situation from occurring, DBMaker allows the database administrator to modify the database backup mode for individual tablespaces when creating them.

To backup BLOB data in a specific tablespace, use the BACKUP BLOB ON option when executing the CREATE TABLESPACE command. To avoid backing up BLOB data in a specific tablespace, use the BACKUP BLOB OFF option when executing the CREATE TABLESPACE command.

The backup mode of each tablespace will then depend on the combination of database backup mode and tablespace backup mode as follows:

- If the database is running in BACKUP-DATA-AND-BLOB mode and a tablespace was created with the BACKUP BLOB ON option, DBMaker will backup BLOB data in that tablespace.
- If the database is running in BACKUP-DATA-AND-BLOB mode and a tablespace was created with the BACKUP BLOB OFF option, DBMaker will not backup BLOB data in that tablespace.
- If the database is running in BACKUP-DATA mode, DBMaker will not backup BLOB data regardless of whether a tablespace was created with the BACKUP BLOB ON or BACKUP BLOB OFF option.

DBMaker uses the BACKUP BLOB ON mode by default for newly created tablespaces. All changes to BLOB data in that tablespace will be recorded in the journal file when the database is in BACKUP-DATA-AND-BLOB mode.

Backup File Object Mode

In addition to backing up regular and BLOB data in the database, users may choose to back up file objects. File objects are backed up only during automatic full backups initiated by the backup daemon. Users should first start the backup server, set the full backup schedule, and set the backup directory. For more information full backup settings refer to section 14.6, *Backup Server*.

There are two types of file objects: user file objects and system file objects. The database administrator may choose to back up user file objects, system and user file objects, or neither. The **dmconfig.ini** keyword **DB_BkFoM** specifies the Backup Mode of File Objects.

- **DB_BkFoM** = 0: Do not backup file objects.
- **DB_BkFoM** = 1: Backup system file objects only.
- **DB_BkFoM** = 2: Backup both system and user file objects.

When backing up file objects (**DB_BkFoM** = 1, 2), the backup server copies all external files of file objects to the "*fo*" subdirectory under the directory specified by

DB_BkDir keyword. The schedule follows the full backup schedule specified by **DB_FBkTm** and **DB_FBkTv** .

➲ **Example**

An excerpt from a **dmconfig.ini** file containing related keywords:

```
[MyDB]
DB_BkSvr = 1                        ; starts the backup server
DB_FBKTm = 01/05/01 00:00:00        ; begins from midnight at May 1, 2001.
DB_FBKTV = 1-00:00:00               ; interval is every one day.
DB_BkDir = /home/dbmaker/backup     ; backup directory
DB_BkFoM = 2                        ; backup both system and user file objects
```

Since the Backup Mode of File Objects is 2, the backup server will copy all external database file objects to the "**/home/dbmaker/backup/FO**" directory. If the **FO** subdirectory does not exist, the backup server will create it.

The files in *FO* subdirectory are renamed with a sequential number. For example, if the name of the original external file is "**/DBMaker/mydb/fo/ZZ000123.bmp**", the backup server would copy it to the **FO** subdirectory and rename it '**fo0000000344.bak**', meaning it is the 344th file object. The mapping between the source full file name and its new name is recorded in the file object mapping list file, **dmFoMap.his**. For more information about the file object mapping list file, refer to section 14.7, *Backup History Files*

The backup server will also move the previous version of file objects to the **FO** subdirectory under the old backup directory specified by **DB_BkOdr**.

Database administrators should consider that enabling file object backup requires more time for a full backup. The cost of complete full backup includes (1) copying the previous full backup if **DB_BkOdr** is set; (2) copying all database files; (3) copying all journal files; and (4) copying all file objects if **DB_BkFoM** is set. Also, ensure that there is enough disk space in the backup directory specified by **DB_BkDir** for all backup files to avoid backup failure.

Setting the Backup Mode

DBMaker provides several different methods to set the backup mode. The method you choose depends on whether your database is online or offline, and whether you are more comfortable editing the configuration file directly, using the dmSQL command line utility, or using the JServer Manager graphical utility.

Modifying the backup mode of a database to provide a higher level of backup protection (i.e. from NONBACKUP to BACKUP-DATA mode, or from BACKUP-DATA to BACKUP-DATA-AND-BLOB mode) has an effect on journal usage. The journal begins recording changes to data that was previously not recorded before modifying the backup mode. As a result, it is necessary to perform a full backup when you change the backup mode. This provides a starting point for the backup journal files to update during the restoration process.

No extra steps are required when modifying the backup mode of a database to provide a lower level of backup protection (i.e. from BACKUP-DATA or BACKUP-DATA-AND-BLOB mode to NONBACKUP mode) since the journal simply stops recording changes to data. DBMaker will use the previous full backup as a starting point for the backup journal files to update during the restoration process. However, some changes to data may be lost if the database is restored after changing to a lower level of backup protection.

The database administrator may change the backup mode of the database while offline using the **dmconfig.ini** file or JServer Manager. Since the backup mode affects journal usage, an offline full backup must be performed before starting the database with the new backup mode setting. Backup modes may be changed from one mode to another without restriction when offline, providing a full backup is made when going from a lower level of backup protection to a higher level. For more information on performing an offline full backup, see "Offline Full Backups" later in this chapter.

A database adminstrator can change the backup mode of a database online using dmSQL. Since the backup mode will affect journal usage, backup mode must be changed from a lower level of backup protection to a higher level (i.e. from NONBACKUP to BACKUP-DATA or BACKUP-DATA-AND-BLOB mode, or from BACKUP-DATA to BACKUP-DATA-AND-BLOB mode) between the start and finish of a full backup period.

➲ **Example**

To use dmSQL to change the backup mode online:

```
dmSQL> begin backup;
dmSQL> set data backup on;
dmSQL> end backup datafile;
dmSQL> end backup journal;
```

or:

```
dmSQL> begin backup;
dmSQL> end backup datafile;
dmSQL> set data backup on;
dmSQL> end backup journal;
```

DBMaker does not allow the database to go from a higher level of backup protection to a lower level unless it is changed to NONBACKUP mode first. To change from BACKUP-DATA-AND-BLOB to BACKUP-DATA mode, first change to NONBACKUP mode and then follow the rules above for changing from a lower level of backup protection to a higher level. The backup mode may be changed from BACKUP-DATA-AND-BLOB or BACKUP-DATA to NONBACKUP at any time; it does not need to be done between the start and finish of a full backup period.

Using the dmconfig.ini Configuration File

If the database is offline, change the backup mode directly using the **DB_BMode** keyword in the **dmconfig.ini** file. The next time the database is started, the new backup mode will be used. If the database is online, changing the value of the **DB_BMode** keyword will have no effect until the database is shut down and restarted. Remember to perform an offline full backup if the backup mode is going to be changed from NONBACKUP to BACKUP-DATA or BACKUP-DATA-AND-BLOB mode or from BACKUP-DATA to BACKUP-DATA-AND-BLOB mode.

➲ **To set the backup mode using the dmconfig.ini file:**

1. Open the **dmconfig.ini** file on the database server using any **ASCII** text editor.

2. Locate the database configuration section for the database.

3. Change the value of the **DB_BMode** keyword to one of the following values:

```
0 - NONBACKUP mode

1 - BACKUP-DATA mode

2 - BACKUP-DATA-AND-BLOB mode
```

4. Restart the database to begin using the new backup mode.

If the **DB_BMode** keyword is not present in the database configuration section for the database, you will have to add the **DB_BMode** keyword to that database configuration section. You can add the keyword on a separate line anywhere between the start of the database configuration section and the start of the next configuration section; the order the keywords appear in is not important. If you do not specify a value for **DB_BMode**, the default value of 0 (NONBACKUP mode) will be used.

USING DMSQL

If the database is online and you are comfortable using the dmSQL command line utility, you can change the backup mode using the SQL SET command. You must execute this command during an online full backup. The new backup mode will be enabled as soon as the command is executed.

➲ **To set the backup mode using the dmSQL command line utility**

1. Connect to the database to change the backup mode using dmSQL.

2. Begin an online full backup using the **BEGIN BACKUP** command.

3. Change the backup mode during the full backup period by issuing one of the following **SET** commands at the dmSQL command prompt:

```
dmSQL> set backup off;

dmSQL> set data backup on;

dmSQL> set blob backup on;
```

4. Complete the online full backup.

The SET BACKUP OFF command corresponds to NONBACKUP mode, the SET DATA BACKUP ON corresponds to BACKUP-DATA mode, and the SET BLOB BACKUP ON command corresponds to BACKUP-DATA-AND-BLOB mode.

USING JSERVER MANAGER

If the database is offline, you can change the backup mode using the JServer Manager graphical utility. JServer Manager will automatically change the value of the **DB_BMode** keyword in the **dmconfig.ini** file. The next time you start the database, the new backup mode will be used. If the database is online, changing the value of the **DB_BMode** keyword will have no effect until the database is shut down and restarted. You must remember to perform an offline full backup if you are going from NONBACKUP to BACKUP-DATA or BACKUP-DATA-AND-BLOB mode or from BACKUP-DATA to BACKUP-DATA-AND-BLOB mode. For directions on how to set the backup mode offline using the JServer Manager graphical utility refer to the *JServer Manager User's Guide*.

14.5 Offline Full Backups

Offline full backups use operating system commands to back up the database. DBMaker provides this option, however, backup server is recommended. Offline full backups necessitate the database be shut down, furthermore, managing the backup sequence is a more complex process.

To perform an offline full backup, you must have read permission on the database files from the operating system, and write permission on the backup directory from the operating system. If you have to shut down the database first, you must have DBA or SYSADM security privileges.

You can perform an offline full backup regardless of the backup mode; the database may be running in NON-BACKUP, BACKUP-DATA, or BACKUP-DATA-AND-BLOB mode. Using an offline full backup, you can restore the database to the point in time the database was shut down.

Note that offline full backup using JServer Manager does not back up file objects. File objects must be copied manually. Be sure to exactly replicate the file and directory structure if restoring a database from an offline full backup. For directions on how to perform an offline full backup using JServer Manager, refer to the *JServer Manager User's Guide*.

OFFLINE FULL BACKUP USING DMSQL

➲ **To perform an offline full backup using dmSQL:**

1. Notify all users that the database will be shut down at a specified time and ask them to disconnect before that time.
2. If the database is running, shut it down using the TERMINATE DB command. If there are any errors while shutting down the database, restart the database, correct the problem, and shut it down again.
3. Examine the **dmconfig.ini** file and list all relevant files and directories, including the file object directory, which require backup.
4. Use operating system commands or utilities to copy the database files, including data files, journal files, file objects, and the **dmconfig.ini** file, to the backup directory or backup device.

14.6 Backup Server

Although DBMaker provides methods for backing up a database manually, you must still remember to perform the backups on a regular basis. To solve this problem, DBMaker provides a convenient and easy way to perform fully automated online full and incremental backups using Backup Server.

The database administrator may also perform backups during runtime with the JServer Manager utility 'Backup by Backup Server'

Backup Server runs in the background and performs online full and incremental backups on a regular schedule, as journal files become full, or both. This flexibility is possible because Backup Server and the database server communicate to determine when a backup should occur, the type of incremental backup to perform, and which backup options to change. Backup Server starts at the same time as the database server, and continues running until you either stop it or shut down the database server. The backup server is

When performing full backups, Backup Server will copy the last full backup from the backup directory to the old directory. Then, it will copy all database files including journal files and **dmconfig.ini** to the backup directory, over writing the previous full backup.

When performing incremental backups, Backup Server will copy necessary journal files to the backup directory.

There are several options used to configure Backup Server. These options control the filename format of the backup files, the location of the backup directory, the location of the old directory, the schedule Backup Server uses to perform backups, the amount a journal file must fill before Backup Server performs an incremental backup, and the way Backup Server saves backup files.

Backup Server also allows backup-related configuration settings to be made during the run time with the dmSQL SetSystemOption stored procedure.

Starting Backup Server

DBMaker will only start Backup Server if your database is in BACKUP-DATA or BACKUP-DATA-AND-BLOB mode and the value of the **DB_BkSvr** keyword is set to 1 in the **dmconfig.ini** file. Backup Server will not start if your database is in NONBACKUP mode or if the value of the **DB_BkSvr** keyword is set to 0.

You do not have to explicitly start Backup Server after setting the **DB_BkSvr** keyword, since DBMaker will automatically start Backup Server while starting the database. Backup Server is disabled by default.

STARTING BACKUP SERVER USING DMCONFIG.INI

If the database is offline, you can enable Backup Server directly using the **DB_BkSvr** keyword in the **dmconfig.ini** file. The next time you start the database, Backup Server will also start. If the database is online, changing the value of the **DB_BkSvr** keyword will have no effect until the database is shut down and restarted.

➲ **To start Backup Server using the dmconfig.ini file:**

1. Open the **dmconfig.ini** file on the database server using any ASCII text editor.

2. Locate the database configuration section for a database to enable Backup Server.

3. Ensure the backup mode of the database is either BACKUP-DATA or BACKUP-DATA-AND-BLOB mode. The database is in BACKUP-DATA mode if the value of **DB_BMode** is set to 1, and it is in BACKUP-DATA-AND-BLOB mode if the value of **DB_BMode** is set to 2.

4. Change the value of the **DB_BkSvr** keyword to 1 to enable Backup Server.

5. Restart the database to begin using Backup Server.

STARTING BACKUP SERVER USING JSERVER MANAGER

If the database is offline, you can enable Backup Server using the JServer Manager graphical utility. JServer Manager will automatically change the value of the **DB_BkSvr** keyword in the **dmconfig.ini** configuration file. The next time you start the database, Backup Server will also start. If the database is online, enabling Backup Server will have no effect until the database is shut down and restarted. For directions on how to start Backup Server while offline using JServer Manager, refer to the *JServer Manager User's Guide*.

Incremental Backup Filename Format

The backup filename format allows you to specify the format Backup Server will use to name incremental backup files. The backup filename format may include both text constants and must include format sequences (escape sequences) that represent special character strings.

An incremental backup file name must consist of at least three special character strings: the full backup id, the database name, and the backup identification number. Backup Server assigns a full backup ID when naming incremental files in a backup sequence. When restoring a database, DBMaker uses the full backup ID to correctly recreate the backup sequence. The database name correctly identifies the database to which an incremental backup file belongs. The backup identification number identifies the relative position of the incremental backup file in the backup sequence.

Format sequences have three parts: the escape character, the length value, and the format character. Valid format sequences are:

%[*x*]F—The full backup ID. The variable *x* may have values 1-4 where the values represent the following formats;

1: full backup id shown as YYYYMMDD, e.g. 20010917

2: full backup id shown as MMDD, e.g. 0917

3: full backup id shown as MMDDhhmm, e.g. 09171305

4: full backup id shown as DDhhmmss, e.g. 17130558

%[*n*]B—The backup identification number.

%[*n*]N—The name of the database the journal file belongs to.

The escape character identifies the start of the format sequence, and is represented by the % symbol. If you want to include the % symbol as a text constant in the backup filename format, you must use two % symbols together (i.e. %%). A single digit or one of the valid format characters shown above must immediately follow the % symbol. If any other characters follow the % symbol the backup filename format is invalid, and DBMaker will return an error.

The length value *n* is an integer value between one and nine that determines the length of the character string generated by the format sequence. If the format sequence returns a string that can be represented in fewer characters than the length value provides then zeros will be appended to it. The database name has zeroes added to the right of the name, while all other values have zeroes added to the left. If the format sequence returns a string that requires more characters than the length value provides, it will be truncated. The database name is truncated from the right, while all other values are truncated from the left. The square brackets enclosing the length value indicate the length value is optional; do not include the square brackets when entering the format sequence. If you do not provide a value for the length, Backup Server will use the full length of the character string generated by the format sequence.

The format character identifies the type of special character string the format sequence will return. The format character must be F, B, or N; using any other character will result in an invalid backup filename format, and DBMaker will return an error. A valid format character that does not immediately follow either the escape character or the escape character and a single digit will be treated as a text constant.

Date and time values are taken from the system. These values will only be correct if the system date and time are correct. The value for the backup identification number is the ordinal position of the backup journal file in the backup sequence. DBMaker automatically provides this number for each journal file that is backed up by Backup Server.

The backup filename format is specified by the **DB_BkFrm** keyword in the **dmconfig.ini** file. If you do not specify a backup filename format, Backup Server will use the default format: %2F%4N%4B.JNL. The total length of the filename returned by the backup filename format must not exceed 256 characters in length.

DBMaker provides several different methods to set the backup filename format. The method you choose depends on whether you are more comfortable editing the configuration file directly or using the JServer Manager graphical utility.

Using dmconfig.ini To Set Backup File Name Format

If the database is offline, you can set the backup filename format used by Backup Server directly using the **DB_BkFrm** keyword in the **dmconfig.ini** file. The next time you start the database, Backup Server will apply this backup filename format to all backup journal files. If the database is online, changing the value of the **DB_BkDir** keyword will have no effect until the database is shut down and restarted.

➲ **To set the backup file format using the dmconfig.ini configuration file:**

1. Open the **dmconfig.ini** file on the database server using any **ASCII** text editor
2. Locate the database configuration section for a database.
3. Change the value of the **DB_BkFrm** keyword to a string containing the format to use for the backup filename format.

NOTE *The string may contain any valid format sequences and text constants, but the total length of the resulting filename must not exceed 256 characters in length.*

4. Restart the database to begin using the new backup filename format.

SETTING BACKUP FILE NAME FORMAT WITH JSERVER MANAGER

If the database is offline or online, you can set the backup filename format used by Backup Server using the JServer Manager graphical utility. JServer Manager will automatically change the value of the **DB_BkFrm** keyword in the **dmconfig.ini** file. The next time you start the database, Backup Server will apply this backup filename format to all backup journal files. For directions on how to set the backup file format using JServer Manager, refer to the *JServer Manager User's Guide*..

Backup Directory

The backup directory specifies where the Backup Server will place backup files. The backup directory must be a directory that already exists. If the directory you wish to use does not exist, you must create it using operating system commands before you specify it as the backup directory. You should choose a backup directory on a different disk than the database files to prevent the loss of both the database and the backup files in the event of a media error.

The backup directory is specified by the **DB_BkDir** keyword in the **dmconfig.ini** file. The value of the **DB_BkDir** keyword may contain either a full or a relative path to the backup directory. If you do not specify a backup directory, the Backup Server will automatically create a default backup directory named **backup** under the database directory. The database directory is specified by the **DB_DbDir** keyword in the **dmconfig.ini** file. The total length of the backup directory path must not exceed 79 characters in length.

It is not a good idea to allow the Backup Server to create and use the default backup directory if you have more than one database in the same directory. In this case, the backup history information from one database may overwrite or append to the backup history information from another database, rendering one or both of the backups unusable. To avoid this type of problem you can put each database in a different database directory, or explicitly specify a backup directory for each database. Placing each database in a different database directory is the preferred method, since this allows you to see exactly which files belong to which database.

DBMaker provides several different methods to set the backup directory. The method you choose depends on whether your database is online or offline, and whether you are more comfortable editing the configuration file directly or using the JServer Manager graphical utility.

Using dmconfig.ini to Set Backup Directory

If the database is offline, you can set the backup directory used by Backup Server directly using the **DB_BkDir** keyword in the **dmconfig.ini** file. The next time you start the database, Backup Server will use this directory as the backup directory. If the database is online, changing the value of the **DB_BkDir** keyword will have no effect until the database is shut down and restarted.

➲ **To set the backup directory using the dmconfig.ini configuration file:**

1. Open the **dmconfig.ini** file on the database server using any ASCII text editor.
2. Locate the database configuration section for a database.
3. Change the value of the **DB_BkDir** keyword to a string containing the name of an existing directory to set the backup directory.
4. Restart the database to begin using the new backup directory.

Using dmSQL To Set Backup Directory On Line

The SetSystemOption command can be used to change the backup directory while the database is running. The general syntax for the command is:

```
SetSystemOption('bkdir', 'path')
```

Where *path* is the full path of the new backup directory. The length of the string in *path* should not exceed 256 characters.

➲ **Example**

To change the directory path to E:/storage/database/backup/WebDB, enter the following line at the dmSQL command prompt.

```
dmSQL> SetSystemOption('bkdir', 'E:/storage/database/backup/WebDB');
```

USING JSERVER MANAGER TO SET BACKUP DIRECTORY

If the database is offline, you can set the backup directory used by Backup Server using the JServer Manager graphical utility. JServer Manager will automatically change the value of the **DB_BkDir** keyword in the **dmconfig.ini** file. The next time you start the database, Backup Server will use this directory as the backup directory. If the database is online, JServer Manager can change the backup directory immediately or delay the change until the next time you restart the database. In either case, JServer Manager will also make a copy of the backup history file in the new backup directory. For directions on how to set the backup directory using JServer Manager, refer to the *JServer Manager User's Guide*.

Setting Multiple Backup Paths

DBMaker also supports multiple backup file paths for users. This function is useful when a user tries to save to a backup path, but the backup path does not provide enough space for the backup to be completed. If the multiple backup option is set DBMaker will then shunt the remaining data to be backed up to secondary backup locations so that the backup can be properly performed. Users are able to use multiple backup paths on full or incremental backups. DBMaker has the following constraints when backing up information using multiple backup paths:

- When a database system attempts to backup files, it will try to store files in the paths one by one for each file. For example, when storing a file to backup directory 1 and the directory does not have enough space to store the file, then the file is shunted to backup directory 2, and so on. If all backup directories are full an error message will be returned.
- Only one backup directory can be used to backup files on the slave sites.
- FOs must backup in the first backup directory.
- The maximum number of backup paths is 32.

➲ **Example:**

When setting multiple backup paths DBMaker conforms to the following structure:


```
DB_BKDIR = <BKDIR 1> <SIZE 1> < BKDIR 2> <SIZE 2> < BKDIR 3> <SIZE 3>…

< BKDIR n > : the n's backup path
< SIZE n > : the size of the n's backup path ( 4k per unit )
```

So when setting multiple backup paths for the database **DB1** you need to set the paths in DB_BKDIR.

```
DB_BKDIR = /home/usr/dbmaker/bk 5000 /home2/backup 1000
```

When the available space in home/usr/dbmaker/bk is full the database will backup at home2/backup.

Setting the Old Directory

The old directory specifies where the Backup Server will place the previous full backup files. You should choose it on a different disk than the database files to prevent the loss of both the database and the backup files in the event of a media error.

The old directory is specified by the **DB_BkOdr** keyword in the **dmconfig.ini** file. If you do not specify it, the Backup Server will not backup the previous full backup files.

USING DMCONFIG.INI TO SET THE OLD DIRECTORY

You can set the old directory used by Backup Server directly using the **DB_BkOdr** keyword in the **dmconfig.ini** file. The next time you start the database, Backup Server will use this directory as the old directory. If the database is online, changing the value of the **DB_BkOdr** keyword will have no effect until the database is shut down and restarted.

USING JSERVER MANAGER TO SET THE OLD DIRECTORY

If the database is offline, you can set the location for the previous backup using the JServer Manager graphical utility. JServer Manager will automatically change the value of the **DB_BkODr** keyword in the **dmconfig.ini** file. The next time you start the database, Backup Server will use this directory as the backup directory. If the database is online, JServer Manager can change the old backup directory immediately or delay the change until the next time you restart the database. For directions on how to set the old backup directory while offline using JServer Manager, refer to the *JServer Manager User's Guide*.

Incremental Backup Settings

The incremental backup schedule specifies the times when Backup Server will perform an online incremental backup. The schedule is composed of two parts: the initial backup time and the interval time. The initial backup time determines the date and time Backup Server will perform the first incremental backup, and the interval time determines the length of time to wait between subsequent incremental backups.

You can combine the incremental backup schedule with the journal trigger value to backup your database both on a regular schedule and when journal files fill to a specified percentage. If you do not specify an incremental backup schedule, Backup Server will not backup the database on a regular schedule. However, Backup Server will continue to perform incremental backs as journal files fill even without an incremental backup schedule.

The initial backup time is specified by the **DB_BkTim** keyword in the **dmconfig.ini** file. You must enter the value of the **DB_BkTim** keyword as a date and time in the format YY/MM/DD HH:MM:SS. There is no default value for the initial backup time. However, if you use JServer Manager to enable Backup Server, JServer Manager will provide a default value for you and write this value into the **dmconfig.ini** file.

The interval time is specified by the **DB_BkItv** keyword in the **dmconfig.ini** file. You must enter the value of the **DB_BkItv** keyword as a time interval in the format D-HH:MM:SS. There is no default value for the interval time. However, if you use JServer Manager to enable Backup Server, JServer Manager will provide a default value of 1-00:00:00 for you and write this value into the **dmconfig.ini** file.

DBMaker provides several different methods to set the incremental backup schedule. The method you choose depends on whether your database is online or offline, and whether you are more comfortable editing the configuration file directly or using the JServer Manager graphical utility.

USING DMCONFIG.INI TO CHANGE INCREMENTAL BACKUP SETTINGS

If the database is offline, you can set the backup schedule used by Backup Server directly using the **DB_BkTim** and **DB_BkItv** keywords in the **dmconfig.ini** configuration file. The next time you start the database, Backup Server will use these settings for the incremental backup schedule. If the database is online, changing the value of the **DB_BkTim** and **DB_BkItv** keywords will have no effect until the database is shut down and restarted.

➲ **To set the backup schedule using the dmconfig.ini file:**

1. Open the **dmconfig.ini** file on the database server using any ASCII text editor.
2. Locate the database configuration section for a database to change the backup schedule.
3. Change the value of the **DB_BkTim** keyword to a date and time using the YY/MM/DD HH:MM:SS value format.
4. Change the value of the **DB_BkItv** keyword to a time interval using the DDDDD-HH:MM:SS value format.
5. Restart the database to begin using the new backup schedule.

USING DMSQL TO CHANGE INCREMENTAL BACKUP SETTINGS

The SetSystemOption command can be used to change the incremental backup start time and interval while the database is running. The general syntax to change the incremental backup start time is:

```
SetSystemOption('bktim', 'StartTime')
```

The general syntax to change the incremental backup interval is:

```
SetSystemOption('bkitv', 'Interval')
```

StartTime is the time to start the first incremental backup, and has the format YY:MM:DD HH:MM:SS. *Interval* is the time interval that incremental backups occur, and has the format D-HH:MM:SS.

➲ **Example**

To set the incremental backup interval to 1 hour, enter the following line at the dmSQL command prompt.

```
dmSQL> SetSystemOption('bkitv', '0-1:00:00');
```

USING JSERVER MANAGER TO CHANGE INCREMENTAL BACKUP SETTINGS

If the database is offline, you can set the incremental backup schedule used by Backup Server using the JServer Manager graphical utility. JServer Manager will automatically change the value of the **DB_BkTim** and **DB_BkItv** keywords in the **dmconfig.ini** file. The next time you start the database, Backup Server will use these settings as the new incremental backup schedule. If the database is online, JServer Manager can change the backup schedule immediately or delay the change until the next time you restart the database. For directions on how to set the incremental backup schedule using JServer Manager, refer to the *JServer Manager User's Guide*.

Journal Trigger Value Settings

The journal trigger value specifies the percentage a journal file must fill before Backup Server will perform an online incremental backup. You can combine the journal trigger value with the backup schedule to backup your database on a regular schedule and when journal files fill to the specified percentage.

The journal trigger value is specified by the **DB_BkFul** keyword in the **dmconfig.ini** file. The value of the **DB_BkFul** keyword may be an integer value in the range 50-100, or zero. Values between 50-100 represent the percentage a journal file must fill before Backup Server performs a backup. A value of zero causes Backup Server to perform a backup whenever a journal file fills completely. Setting the value to 0 is effectively the same as setting it to a value of 100, since both will cause Backup Server to perform a backup whenever a journal file fills completely (100% full). If you do not specify a value for the journal trigger value, Backup Server will use the default value of zero.

DBMaker provides several different methods to set the journal trigger value. The method you choose depends on whether your database is online or offline, and whether you are more comfortable editing the configuration file directly or using the JServer Manager graphical utility.

Using dmconfig.ini To Change The Journal Trigger Value

If the database is offline, you can set the journal trigger value used by Backup Server directly using the **DB_BkFul** keyword in the **dmconfig.ini** file. The next time you start the database, Backup Server will use this setting for the journal trigger value. If the database is online, changing the value of the **DB_BkFul** keyword will have no effect until the database is shut down and restarted.

➲ **To set the journal trigger value using the dmconfig.ini file:**

1. Open the **dmconfig.ini** file on the database server using any ASCII text editor.
2. Locate the database configuration section for a database to change the journal trigger value.
3. Change the value of the **DB_BkFul** keyword to an integer value between 50-100, or set it to zero.
4. Restart the database to begin using the new journal trigger value.

Using dmSQL To Change The Journal Trigger Value

The SetSystemOption command can be used to change the journal trigger value while the database is running. The general syntax to change the incremental backup start time is:

```
SetSystemOption('bkful', 'n')
```

Where *n* is either 0 or 50-100. Setting *n* to 0 will trigger the backup server whenever a journal file is full. Setting *n* between 50-100 specifies the percentage a journal file fills to before the backup server activates.

➲ **Example**

To set the journal trigger value to 75 percent, enter the following line at the dmSQL command prompt.

```
dmSQL> SetSystemOption('bkful', '75');
```

USING JSERVER MANAGER TO CHANGE THE JOURNAL TRIGGER VALUE

If the database is offline, you can set the journal trigger value used by Backup Server using the JServer Manager graphical utility. JServer Manager will automatically change the value of the **DB_BkFul** keyword in the **dmconfig.ini** file. The next time you start the database, Backup Server will use this setting as the new journal trigger value. If the database is online, JServer Manager can change the journal trigger value immediately or delay the change until the next time you restart the database. For directions on how to set the journal trigger value using JServer Manager, refer to the *JServer Manager User's Guide*.

Compact Backup Mode Settings

Compact backup mode specifies whether Backup Server will backup entire journal files or only full journal blocks when it performs an online incremental backup. This is possible since not every journal block contains data needed to restore a database, so Backup Server will only backup the necessary journal blocks when it performs a backup. This allows you to save storage space on your backup device, but it also means restoring a database may take more time.

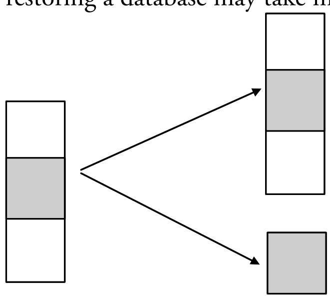

Non-Compact Mode:
Backup entire Journal files

Compact Mode:
Backup necessary Journal blocks

The compact backup mode setting is specified by the **DB_BkCmp** keyword in the **dmconfig.ini** configuration file. The value of the **DB_BkCmp** keyword may be zero or one. Setting the value to one enables compact backup mode, and setting it to zero disables compact backup mode. If you do not specify a value for the compact backup mode, Backup Server will use the default value of one (enabled).

DBMaker provides several different methods to set the compact backup mode. The method you choose depends on whether your database is online or offline, and whether you are more comfortable editing the configuration file directly or using the JServer Manager graphical utility.

USING DMCONFIG.INI TO SET COMPACT BACKUP MODE

If the database is offline, you can set the compact backup mode setting used by Backup Server directly using the **DB_BkCmp** keyword in the **dmconfig.ini** file. The next time you start the database, Backup Server will use this setting for the compact backup mode. If the database is online, changing the value of the **DB_BkCmp** keyword will have no effect until the database is shut down and restarted.

➲ **To set the Compact Backup Mode using the dmconfig.ini configuration file:**

1. Open the **dmconfig.ini** file on the database server using any ASCII text editor.
2. Locate the database configuration section for a database to change the journal trigger value.
3. Change the value of the **DB_BkCmp** keyword to one to enable compact backup mode, or zero to disable compact backup mode.
4. Restart the database to begin using the new journal trigger value.

USING JSERVER MANAGER TO SET COMPACT BACKUP MODE

If the database is offline, you can set the compact backup mode setting used by Backup Server using the JServer Manager graphical utility. JServer Manager will automatically change the value of the **DB_BkCmp** keyword in the **dmconfig.ini** file. The next time you start the database, Backup Server will use this setting as the new compact backup mode setting. If the database is online, JServer Manager can change compact backup mode setting immediately or delay the change until the next time you restart the database. For directions on how to set the Compact Backup Mode using JServer Manager, refer to the *JServer Manager User's Guide*.

Full Backup Schedule

The full backup schedule specifies the times when Backup Server will perform an online full backup. The schedule is composed of two parts: the initial backup time and the interval time. The initial backup time determines the date and time Backup Server will perform the first full backup, and the interval time determines the length of time to wait between subsequent full backups.

You can combine the full backup schedule with an incremental to backup your database. If you do not specify an full backup schedule, Backup Server will not perform full backups on a regular schedule.

The initial backup time is specified by the **DB_FBkTm** keyword in the **dmconfig.ini** file. You must enter the value of the **DB_FBkTm** keyword as a date and time in the format YY/MM/DD HH:MM:SS. There is no default value for the initial backup time.

The interval time is specified by the **DB_FBkTv** keyword in the **dmconfig.ini** file. Enter the value of the **DB_FBkTv** keyword as a time interval in the format D-HH:MM:SS. There is no default value for the interval time.

USING DMCONFIG.INI TO SET THE FULL BACKUP MODE

If the database is offline, you can set the full backup schedule used by Backup Server directly using the **DB_FBkTm** and **DB_FBkTv** keywords in the **dmconfig.ini** file. The next time you start the database, Backup Server will use these settings for the full backup schedule. If the database is online, changing the value of the **DB_FBkTm** and **DB_FBkTv** keywords will have no effect until the database is shut down and restarted.

➲ **To set the full backup schedule using the dmconfig.ini configuration file:**

1. Open the **dmconfig.ini** file on the database server using any ASCII text editor.
2. Locate the database configuration section for a database to change the journal trigger value.

3. Set the configuration parameter **DB_FBkTm** to a value of the format YY/MM/DD HH:MM:SS, and **DB_FBkTv** to a value of the format D-HH:MM:SS.

4. Restart the database to begin using the new full backup schedule.

Using JServer Manager To Set the Full Backup Schedule

You can set the full backup schedule with JServer Manager by using the start database setup utility. JServer Manager will automatically change the value of the **DB_FBkTm** and **DB_FBkTv** keywords in the **dmconfig.ini** file. The next time you start the database, Backup Server will use this setting as the new full backup schedule. For directions on how to set the full backup schedule using JServer Manager, refer to the *JServer Manager User's Guide*.

Backup Mode of File Objects

The Backup Mode of File Objects lets the database administrator decide whether Backup Server will back up file objects during a full backup. It is also possible to specify Backup Server to back up just system file objects or system and user file objects.

It is possible to set the Backup Mode of File Objects in a number of ways. The configuration keyword **DB_BkFoM** determines the setting during database startup, but it may also be modified during runtime with dmSQL or the JServer Manager utility.

The backup server will move all files from the previous backup to the old backup directory specified by **DB_BkOdr**.

Starting file object backup will cause the database to require more time to complete a full backup, depending on how many file objects are in the database. The total cost of a complete full backup includes (1) copying the previous full backup if **DB_BkOdr** is set; (2) copying all database files; (3) copying all journal files; and (4) copying all file objects if **DB_BkFoM** is set. Be sure that enough disk space is available in the backup directory specified by **DB_BkDir** (and **DB_BkOdr** if applicable) for all mentioned backup files to avoid backup failure.

File objects are copied into an FO directory that is created in the backup directory at the time a full backup is performed. File objects are renamed sequentially when they are copied to the directory for backed-up file objects. The files in the **/fo** subdirectory are renamed starting with the letters FO followed by a ten digit serial number. All backup file objects are appended with the file extension .BAK. The mapping between the source file name and path and the backup file name is recorded in the file object mapping file **dmFoMap.his**.

THE BACKUP FILE OBJECT MAPPING FILE

The file object mapping file **dmFoMap.his** is created in the "**DB_BkDir/FO**" directory. It is a pure ASCII text file that records the original external file name and backup file name. The format looks like:

```
Database Name: MYDB
Begin Backup FO Time: 2001.5.13 2:33
FO Backup Directory: /DBMaker/mydb/backup/FO (i.e. DB_BkDir/FO)
[Mapping List]
s, fo0000000000.bak, "/DBMaker/mydb/fo/ZZ000001.bmp"
u, fo0000000001.bak, "/home2/data/image.jpg"
....
s, fo0000002345.bak, "/DBMaker/mydb/fo/ZZ00AB32.txt"
```

The content before “[Mapping List]” is only a description for user reference. Each line after "[Mapping List]" represents a record that shows the file object type (s = system file object, u = user file object), the new file in **/fo** subdirectory and its original file name and path. This mapping file is necessary for restoration of file objects.

SETTING THE BACKUP MODE OF FILE OBJECTS WITH DMCONFIG.INI

The configuration file keyword **DB_BkFoM** determines the backup mode of file objects:

- **DB_BkFoM** = 0: Do not back up file objects
- **DB_BkFoM** = 1: Back up system file objects only
- **DB_BkFoM** = 2: Back up both system and user file objects

If **DB_BkFoM** = 1 or 2, the backup server will copy all file objects to the **/fo** subdirectory under the backup directory. The schedule follows the full backup schedule.

➲ **Example**

An entry in a **dmconfig.ini** file for specifying the file object backup parameters.

```
[MyDB]
DB_BkSvr = 1                        ; starts the backup server
DB_FBKTm = 01/05/01 00:00:00        ; begins at midnight, May 1, 2001.
DB_FBKTV = 1-00:00:00               ; interval is once every day.
DB_BkDir = /home/dbmaker/backup     ; backup directory
DB_BkFoM = 2                        ; backup both system and user file objects
```

Since the backup mode is 2, the backup server will copy all external files (user file objects) and system file objects to the **/home/dbmaker/backup/FO** directory. If the **FO** subdirectory does not exist, the Backup Server will create it.

Using dmSQL To Set the Backup Mode of File Objects

The SetSystemOption command can be used to change the backup mode of file objects while the database is running. The general syntax to change the Backup Mode of File Objects is:

```
SetSystemOption('bkfom', 'n')
```

Where *n* is 0, 1, or 2. Setting *n* to 0 will turn the Backup Mode of File Objects to off. Setting *n* to 1 configures backup server to back up all system file objects during a full backup. Setting *n* to 2 configures backup server to back up all system and user file objects during a full backup.

➲ **Example**

To configure Backup Server to perform a full backup on all user and system file objects, enter the following line at the dmSQL command prompt.

```
dmSQL> SetSystemOption('bkfom', '2');
```

SETTING THE BACKUP MODE OF FILE OBJECTS WITH JSERVER MANAGER

The settings under the **Backup File Object Mode** effect how file objects are copied during the full backup process. Selecting **Do Not Backup File Objects** disables file backup during the full backup process. Selecting **Backup System File Objects Only** will result in system file objects being backed up during automatic full backups. Selecting **Backup System and User File Objects** will result in both system file objects and user file objects being copied to the backup directory during automatic full backups. For directions on how to set the Backup Mode of File Objects during database startup or with the Run Time Settings dialog in JServer Manager, refer to the *JServer Manager User's Guide*.

Stopping Backup Server

Stop Backup Server when you no longer want it to run by setting the value of the **DB_BkSvr** keyword to zero. However, Backup Server will continue running until you shut down your database and restart it.

STOPPING BACKUP SERVER USING DMCONFIG.INI

If the database is offline, you can disable Backup Server directly using the **DB_BkSvr** keyword in the **dmconfig.ini** file. The next time you start the database, Backup Server will not start. If the database is online, changing the value of the **DB_BkSvr** keyword will have no effect until the database is shut down and restarted.

➲ **To stop Backup Server using the dmconfig.ini file:**

1. Open the **dmconfig.ini** file on the database server using any ASCII text editor.
2. Locate the database configuration section for a database to change the backup mode.
3. Change the *value* of the **DB_BkSvr** keyword to 0 to disable the Backup Server.
4. Restart the database.

STOPPING BACKUP SERVER USING JSERVER MANAGER

If the database is offline, you can disable Backup Server using the JServer Manager graphical utility. JServer Manager will automatically change the value of the **DB_BkSvr** keyword in the **dmconfig.ini** configuration file. The next time you start the database, Backup Server will not start. If the database is online, disabling Backup Server will have no effect until the database is shut down and restarted. For directions on how to stop Backup Server using JServer Manager, refer to the *JServer Manager User's Guide*.

14.7 Backup History Files

Manual incremental backups require you to note which journal files were backed up, when they were backed up, and where the backup files are located. This information may be misplaced or lost if you are not careful. Automatic backups using the backup server prevent this by automatically storing this information in the backup history file.

Locating the Backup History File

This file is created in the backup path and is named ***dmBackup.his***. The file will automatically be used during restoration of a database.

Understanding the Backup History File

Backup history files contain all information pertaining to the id number, file names, and time and date that backups were made. DBMaker uses the backup history file to track backup sequences and ensure the consistency of full and incremental backups within each sequence.

The following is the format of the backup history file:

```
<backup_id>: file_name -> archive_file_name: time: event
```

This denotes that a file named *file_name* was copied to an archive file named *archive_file_name* at time because of event. The event is a text string indicating the reason for the backup. This string can be JOURNAL-FULL, TIME-OUT, ON-LINE-FULL-BACKUP-BEGIN, ON-LINE-FULL-BACKUP, or ON-LINE-FULL-

BACKUP-END. The string JOURNAL-FULL indicates an incremental backup was performed because the journal was full. The string TIME-OUT indicates an incremental backup was performed because the scheduled backup interval has elapsed. The string ON-LINE-FULL-BACKUPxxxx means it is a full backup.

Using the Backup History File

If journal full occurs frequently, lower the backup journal full percentage or shorten the time interval. Also, find out if the backup interval is too short by checking the backup history file. If the same journal file is backed up consecutively in the backup history file, the time interval may be too short. This situation will waste disk space because each file may only contain a few changed blocks. To avoid this, enable compact backup mode or lengthen the backup time interval.

If many journal files are backed up every time, it may mean the time interval is too long. This situation is more dangerous because of the possibility of losing more data when a disk fails. Ideally, one journal file should be backed up every time an incremental backup occurs. This will save storage space and lower the risk of losing journal data.

To shorten the time of recovery from media failures, perform full backups regularly, even if you are using the backup server. In addition, this will also reduce the amount of backup storage needed.

After the backup server is running, you still can perform incremental backups manually. As stated above, note the backup ID, time, and location for the backup files.

Understanding the File Object Backup History File

The file object backup history file, **dmFoMap.his**, keeps a record of all file objects that have been backed up by setting the file object backup configuration parameter on. **dmFoMap.his** is placed in the "**<DB_BkDir>/FO**" directory, is a pure ASCII text file that records the original external file name and backup file name.

The following is the file format:

```
Database Name: MYDB
```

```
Begin Backup FO Time: 2001.5.13 2:33
FO Backup Directory: /DBMaker/mydb/backup/FO (i.e. DB_BkDir/FO)
[Mapping List]
s, fo0000000000.bak, "/DBMaker/mydb/fo/ZZ000001.bmp"
u, fo0000000001.bak, "/home2/data/image.jpg"
....
s, fo0000002345.bak, "/DBMaker/mydb/fo/ZZ00AB32.txt"
```

In the first column, **s** or **u** represent system or user file objects, respectively. The second column gives the backup name, and the third column gives the full name and path of the original file object.

14.8 Recovery Options

Restoring a database will recreate the database, as it existed at the time of the most recent full backup, plus any changes that have been made in the backed up journal files.

Analyzing Recovery Options

What recovery operations are available?

- The answer to this question is determined by whether or not a database is in BACKUP mode. If the database is operating in NONBACKUP mode, the only option for restoration after a disk failure is to restore the most recent full backup and restart the database. All work performed since the last full backup will be lost, and must be re-entered. If this is the case, there is no need to answer the following questions.
- If the database is operated in BACKUP (BACKUP-DATA or BACKUP-DATA-AND-BLOB) mode, several recovery options are available for reconstructing the damaged database.

Preparing for Restoration

Before you restore a database after a disk error, answer the following questions:

- What point in time do you want a database restored to?
- If your answer is the time when the disk error occurs, backup all journal files of the damaged database. These files will help *DBMaker* to restore the database to the most current time.
- What files have previously been backed up?
- Find out where the most current full backup and all subsequent incremental backups are located. For example, suppose you perform a full backup on the 30th day of every month and an incremental backup every 10 days. If your system is damaged on May 25th, you need the full backup from April 30th, the incremental backups from May 10th and May 20th, and the damaged journal files from May 25th. After locating these files, *DBMaker* can restore your database to the state it was in before the failure on May 25. Since all online backup information is stored in the backup history file, *DBMaker* will read it to get this information when restoring a database.

Performing a Restoration

DBMaker uses JServer Manager to perform restoration.

➲ **The following procedure is used to restore a database:**

1. JServer Manager copies all files from the most recent full backup, (*BLOB* files, data files, journal files, and **dmconfig.ini**), to the directory specified by the **DB_DbDir** keyword in the **dmconfig.ini**.
2. To re-create the state the database was in at a specific point in time, set the time to restore the database. Skip this step if restoring to the most current time.
3. Specify the location of all incremental back-up journal files and list the files in order of backup ID numbers.
4. Backed up journal files created after a disk error occurred could be included in the restoration process or after all other backed up journal files.
5. After complete restoration of files and restarting the database in a consistent state, users may begin using the database.

15 Distributed Databases

This chapter introduces the distributed database management functions provided by DBMaker, including distributed databases, the distributed architecture, distributed data access, distributed database object management, and distributed transaction management.

15.1 Introduction to Distributed Databases

Traditional client-server DBMS, as shown in Figure 15-1, locate the database on a specific network computer, and the computer is responsible for handling all client requests.

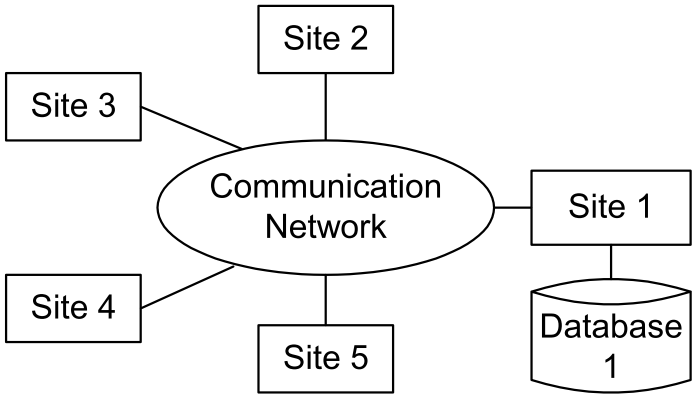

Figure 15-1 Traditional client/server database management system

Distributed databases, as shown in Figure 15-2, locate a copy of the database on several network computers, and each can independently support clients. The distributed database management system manages the databases on these computers, so users can access the data transparently.

DBMaker supports a true distributed architecture to provide a complete and robust distributed database management system (Distributed DBMS). It provides remote database connections, distributed queries, and distributed transaction management. DBMaker also provides table and database replication to keep data automatically up-to-date.

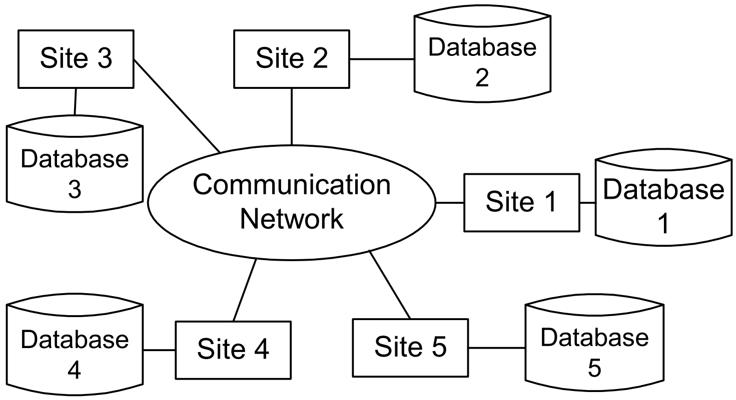

Figure 15-2 Distributed database in client-server

In the DBMaker distributed database environment, you can write application programs using the DBMaker ODBC 3.0 compatible API or perform ad-hoc SQL queries that access data from different parts of the distributed database. DBMaker will transparently integrate the data and return the results, just as if they all came from a local database.

In this chapter, we will briefly describe the system architecture and basic functions of distributed database management using DBMaker. This includes configuring the distributed environment, managing remote data links and distributed transactions, and performing distributed queries. Whether you are a database administrator or an application developer, this chapter will provide you with a thorough overview of the simplicity and power of the DBMaker distributed architecture.

15.2 Distributed Database Structure

The DBMaker distributed database environment builds on the traditional client/server architecture, effectively linking multiple client applications and multiple database servers. Client applications process user requests and display the results, and the database servers handle data management. Each client has a direct connection to a single database server, which is known as the Coordinator Database to that client.

Through the Coordinator Database, the client can connect to other remote databases, which are known as Participant Databases.

DBMaker uses a hierarchical distribution structure to connect to remote databases. This allows DBMaker to access data from a remote database with no direct connection to the Coordinator Database by routing through one of the Participant Databases. When this happens, the Participant Database becomes a Local Coordinator Database, and acts as coordinator for any child databases accessed through it.

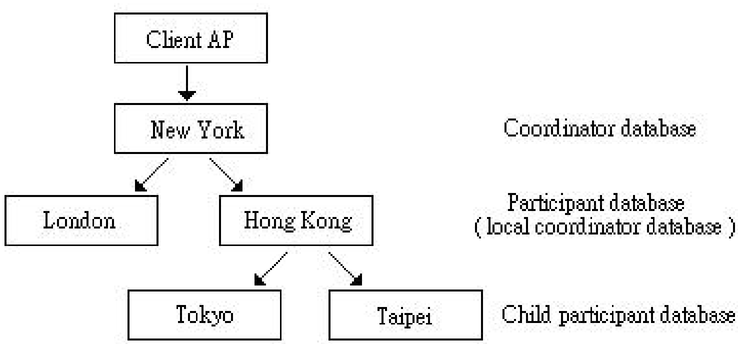

Figure 15-3 DBMaker distribution structure

In Figure 15-3, the client application program connects to the database server in New York, which makes the database in New York the Coordinator Database. If you use the database in New York to access data from London and Hong Kong, then both the London and Hong Kong databases are Participant Databases.

Some of the data you are looking for in Hong Kong might actually be in the databases in Tokyo or Taipei, so the databases in Tokyo and Taipei are Child Participant Databases. This makes the database in Hong Kong a Coordinator Database for the databases in Tokyo and Taipei, so the database in Hong Kong is not only a Participant Database, but also acts as a Local Coordinator Database.

15.3 Distributed Database Environment

Setting up a distributed database environment using DBMaker is very simple. All you need to do is add some keywords to the **dmconfig.ini** file to set the distributed database configuration options. Optionally, these parameters may be set using the JConfiguration Tool. For more information, refer to the *JConfiguration Tool User's Guide*.

You must provide values for the following keywords when setting up a distributed database environment in DBMaker. Keywords with the prefix DB_ are for the client/server connection between the client and the Coordinator Database, and keywords with the prefix DD_ are for the distributed database connections between the Coordinator Database and the Participant Databases.

- **DB_SvAdr=<ip_address/host name>**—specifies the IP address or host name of the Coordinator Database.
- **DB_PtNum=<port number>**—specifies the port number that the client application and the Coordinator Database should use to communicate.
- **DD_DDBMd=<0/1>**—enables distributed database mode for the Coordinator Database. The default value is 0, which means that distributed database mode is disabled.
- **DD_CTimO=<number of seconds>**—specifies the time in seconds that the Coordinator Database should wait when trying to establish a connection to a Participant Database. The default value is 5 seconds.
- **DD_LTimO=<number of seconds>**—specifies the time in seconds that the Coordinator Database should wait when trying to establish a lock on the requested data in a Participant Database. The default value is 5 seconds.
- **DD_GTSVR=<0/1>**—enables the global transaction recovery daemon (GTRECO). The default value is 1, which means the global transaction recovery daemon is enabled.

- **DD_GTItv=<YYYY/MM/DD hh:mm:ss>**—specifies the time interval that the global transaction recovery daemon (GTRECO) should wait when processing pending global transactions.

DBMaker supports an automatic recovery mechanism for distributed transactions that have failed due to network problems or errors on the Participant Database. The automatic recovery mechanism is handled by the GTRECO daemon, which will check whether a distributed database server has any problems with pending global transactions. If any problems are detected, the GTRECO daemon will attempt to recover the pending global transactions. The GTRECO daemon is enabled using the **DD_GTSVR** keyword in the **dmconfig.ini** file.

To better understand how DBMaker manages distributed databases, refer to the following example.

➲ **Example**

The ABC Bank has two branch offices. One branch office is in Los Angeles, and the other one is in Seattle. Each branch maintains their own customer and business data, but the government and financial database is under the central control of the Los Angeles branch.

The LA branch database server **dmconfig.ini** file:

```
[BankTranx]                              ;LA branch business database
DB_DBDIR = c:\database
DB_SVADR = 192.168.0.1
DB_PTNUM = 21000
DD_DDBMD = 1

[BankMIS]                                ;government and financial database
DB_DBDIR = c:\database
DB_SVADR = 192.168.0.1
DB_PTNUM = 30000
DD_DDBMD = 1

[BankTranx@Seattle]                      ;Seattle branch business database
DB_SVADR = 192.168.0.2
DB_PTNUM = 21000
```

```
DD_CTIMO = 20
DD_LTIMO = 10
```

The Seattle branch database server **dmconfig.ini** file:

```
[BankTranx]                          ;Seattle branch business database
DB_DBDIR = c:\database
DB_SVADR = 192.168.0.2
DB_PTNUM = 21000
DD_DDBMD = 1

[BankMIS]                            ;government and financial database
DB_SVADR = 192.168.0.1
DB_PTNUM = 30000
DD_CTIMO = 20

[BankTranx@La]                       ;LA branch business database
DB_SVADR = 192.168.0.1
DB_PTNUM = 21000
DD_CTIMO = 20
DD_LTIMO = 10
```

The LA client application server **dmconfig.ini** file:

```
[BankTranx]                          ;LA branch business database
DB_SVADR = 192.168.0.1
DB_PTNUM = 21000
```

The Seattle client application server **dmconfig.ini** file:

```
[BankTranx]                          ;Seattle branch business database
DB_SVADR = 192.168.0.2
DB_PTNUM = 21000
```

In the files shown above, set **DD_DDBMd** =1 in the configuration section for the local database to enable distributed database support. In these examples, place the keywords in the **BankTranx** configuration section of the Los Angeles and Seattle **dmconfig.ini** files.

In addition, include a database configuration section for the Participant Database in the Coordinator Database configuration file, and for the Coordinator Database in the Participant Database configuration file. In this case, both the Los Angeles branch database and the Seattle branch database use the same database name. If you use the remote database name for the name of the database configuration section, it will cause a conflict with the local database name in the **dmconfig.ini** file.

To avoid this type of problem when using distributed databases, DBMaker can distinguish the remote database name from the local database name by appending a server host description to the remote database name in the local **dmconfig.ini** file.

The remote database name would look like:

```
database_name@server_host_description
```

The server host description can be any identifying name, such as the IP address or host name of the database server, the domain name, or almost any other descriptive text. In this example, the Los Angeles branch client application would use **BankTranx@Seattle** when it wants to access data in the Seattle branch database, and the Seattle branch client application would use **BankTranx@La** when it wants to access data in the Los Angeles branch database.

Also, set up the server address and port name for both the local database and the remote database in their respective configuration sections in the configuration files at both the Los Angeles and Seattle branches.

In this example, the Los Angeles branch configuration file would contain the local server address in the **BankTranx** configuration section, and would contain the Seattle branch server address in the **BankTranx@Seattle** configuration section. Similarly, the Seattle branch configuration file would contain the Los Angeles branch server address in the **BankTranx** configuration section, and would contain the Seattle branch server address in the **BankTranx@La** configuration section.

You should also set the **DD_CTimO** and **DD_LTimO** remote connection parameters. These parameters go in the configuration section for the Participant Database in the Coordinator Database configuration file, and for the Coordinator Database in the Participant Database configuration file.

Every database server in the network can operate on distributed database objects. Any of these database servers can be accessed through the Coordinator Database, in a manner similar to the normal client/server architecture. The SQL commands that reference a remote database will be passed to the remote database server through the Coordinator Database. The Coordinator Database will decompose this SQL command into the local and remote portions, and send the appropriate commands to the remote database server. The Coordinator Database will wait for the remote database to return results, and then merge all local and remote data and return the combined results.

15.4 Distributed Database Objects

DBMaker provides several different methods to access a Participant Database:

- Specify the Participant Database name directly.
- Using database links defined in the Coordinator Database.
- Through remote object mapping such as views or synonyms.

The difference between the first two approaches is that database links contain security information in addition to the remote database name. This allows you to specify the user name and password that you want to use in the database link when you access the remote database.

There is no obvious difference between the statements of a distributed query and a normal query, except in the way database objects are specified. However, when using a remote database, the only remote database objects that can be accessed are tables, views, or synonyms. To access a remote database object, provide the remote database name or database link when specifying the name of the database object. This provides two ways to identify a remote database object:

- remote_database_name:object_owner.object_name
- database_link:object_owner.object_name

➲ **Example 1**

To specify a remote database object in a query:

```
dmSQL> SELECT * FROM Bank:EmpTaple;
dmSQL> DELETE FROM Bank:EmpTable WHERE id = 101;
dmSQL> INSERT INTO Link1:mis.account VALUES (2003,'Kevin Liu','2327-0021');
```

➲ **Example 2**

To access remote database objects in two different Participant Databases:

```
dmSQL> SELECT * FROM ABCBank@La:account a,
                     ABCBankMIS@Seattle:account b
                WHERE a.name = b.name;
```

Remote Database Connections-Using Names

Users can connect to remote databases with the database name of the Coordinator Database Server. Users must know the remote database name, which is defined in the **dmconfig.ini** file in the Coordinator Database Server.

➲ **Example 1**

A client application in the Los Angeles ABC bank branch accesses the database located in the Taipei ABC bank branch. It appears that the user is connecting to the Taipei branch with the user name SYSADM and the password aa. In reality, the user is connecting to the Coordinator Database, which is the Los Angeles branch database. The Coordinator Database then connects to the remote database with the account and password used.

```
dmSQL> CONNECT TO BankTranx SYSADM aa;
dmSQL> SELECT * FROM BankTranx@Taipei:SYSADM.Account ORDER BY AccID;
```

➲ **Example 2**

Using joins to access remote database objects:

```
dmSQL> SELECT * FROM BankMIS:SYSADM.Personnel ORDER BY PID;
dmSQL> SELECT Personnel.* FROM BankTranx@Taipei:Account A,
                               BankMIS:Personnel B
                         WHERE A.CustID = B.CustID;
```

Remote Database Connections-Using Links

A database link creates a connection to a remote database, and contains the login information and password necessary for connecting. The link permits users to connect to a remote database with a different user name than in the Coordinator Database, or to connect to a remote database with no account. It also makes data in a distributed database environment location transparent. The link definition, which also contains the login information and password, is stored in the Coordinator Database.

CREATING DATABASE LINKS

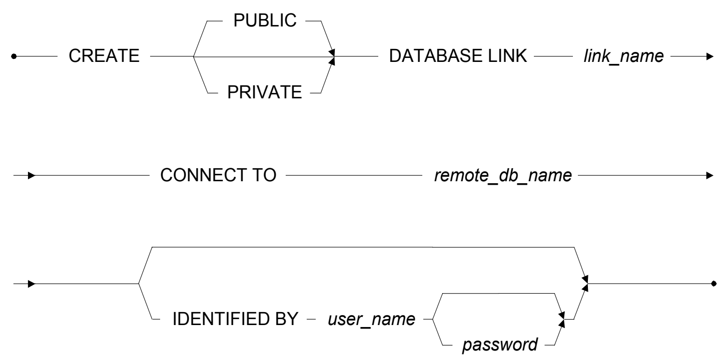

Figure 15-4 Syntax for creating a database link

Only database administrators can create public links to be used by all users in a database. Any user can create private links for themselves. Users may create private links using the same name. A private link will override a public link with an identical name.

DBMaker will create a private link by default if the user does not specify the type of link desired. If the user does not specify the login account and password in the IDENTIFY BY clause, the user's current login name and password will be used by default.

REMOTE DATABASE OBJECTS & DATABASE LINKS

➲ **Example 1**

The following shows how to access remote database objects using database links. In this example, the SYSADM connects to the database and creates a public link named **Bank_Seattle** that connects to the Seattle branch database using the SYSADM account. The SYSADM updates some values and disconnects. Then **user1** connects and performs a query on the **Account** table.

```
dmSQL> CONNECT TO BankTranx SYSADM;
dmSQL> CREATE PUBLIC DATABASE LINK Bank_Seattle CONNECT TO BankTranx@Seattle
    2> IDENTIFIED BY SYSADM;
dmSQL> UPDATE Bank_Seattle:Account SET balance = balance + 100
    2> WHERE id = 1001;
dmSQL> DISCONNECT;
dmSQL> CONNECT TO BankTranx user1 pwd1;
dmSQL> SELECT * FROM Bank_Seattle:Account;
```

➲ **Example 2**

The SYSADM does not specify the account to use when connecting to the public link. This means that when user1 uses the public link to connect to the **Bank_Seattle** database, there must be an account for **user1** in the remote database, and the user1 account must have the authority to query the **SYSADM.Account** table. Otherwise, an error will occur.

```
dmSQL> CONNECT TO BankTranx SYSADM;
dmSQL> CREATE PUBLIC DATABASE LINK Bank_Seattle CONNECT TO BankTranx@Seattle;
dmSQL> SELECT * FROM Bank_Seattle:Account;
dmSQL> DISCONNECT;
dmSQL> CONNECT TO BankTranx user1 pwd1;
dmSQL> SELECT * FROM Bank_Seattle:SYSADM.Account;
```

If a database link name is the same as the remote database name, DBMaker will use the database link name in preference to the remote database name. If you want to access the remote database directly, you must specify the remote database name in the form of dbname@ to force DBMaker to access the remote database directly instead of through the database link.

The following 2 examples show different ways to access a remote database, one through a database link and the other by specifying the remote database name in the form of dbname@"".

➲ **Example 1**

SYSADM connects to a remote database with a link.

```
dmSQL> CONNECT TO BankTranx SYSADM;
```

```
dmSQL> CREATE PUBLIC DATABASE LINK BankMIS CONNECT TO BankMIS
    2> IDENTIFIED BY SYSADM;
dmSQL> DISCONNECT;
```

➲ **Example 2**

user1 connects to a remote database using the **BankMIS@””:SYSADM.Personnel** form.

```
dmSQL> CONNECT TO BankTranx user1 pwd1;
dmSQL> SELECT * FROM BankMIS:Personnel;             //using database link
dmSQL> SELECT * FROM BankMIS@””:SYSADM.Personnel;    //using remote db name
```

DELETING DATABASE LINKS

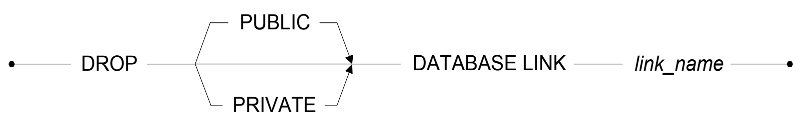

Figure 15-5 Syntax for deleting a database link

Only database administrators can delete public links, and only the owner of a private link can delete it. Ensure to specify the public link to be deleted when it has the same name as a private link or DBMaker will delete the private link by default.

➲ **Example**

To delete a public database link named **BankMIS**:

```
dmSQL> DROP PUBLIC DATABASE LINK BankMIS;
```

Database Object Mapping

Database Object Mapping provides better location transparency in a distributed database environment. There is no difference between the way a user accesses remote database objects with Database Object Mapping and local database objects. This type of Database Object Mapping includes using views and synonyms.

SYNONYMS

Using a synonym to define a remote database object is done by assigning the remote database object an alias name. The privileges you have in the remote database when using a synonym are the same as in a local database.

➲ **Example**

To access a remote database object using a synonym:

```
dmSQL> CONNECT TO BankTranx user1;
dmSQL> CREATE DATABASE LINK LK1 CONNECT TO BankMIS IDENTIFIED BY user2;
dmSQL> CREATE SYNONYM s1 FOR BankTranx:Account;
dmSQL> CREATE SYNONYM s2 FOR LK1:user2.Personnel;
dmSQL> SELECT * FROM s1;
     // SELECT * FROM BankTranx:user1.Account; (BankTranx, user1)
dmSQL> SELECT * FROM s2;
     // SELECT * FROM LK1:user2.Personnel; (BankMIS, user2)
dmSQL> DISCONNECT;
dmSQL> CONNECT TO BankTranx user3;
dmSQL> CREATE DATABASE LINK LK1 CONNECT TO BankMIS IDENTIFIED BY user4;
dmSQL> SELECT * FROM s1;
       // SELECT * FROM BankTranx:user3.Account; (BankTranx, user3)
dmSQL> SELECT * FROM s2;
       // SELECT * FROM LK1:user2.Personnel; (BankMIS, user4)
```

The comments indicate the equivalent SQL expression, the database being connected to, and the account used to connect.

VIEWS

Using a view to define a remote database object is a bit different than using a synonym. The view is not just an alias, but includes the database name, user account, password, object owner, and object name as part of the definition. The privileges a user has in the remote database depend on the privileges of the user that created it.

➲ **Example**

To access a remote database object using a view:

```
dmSQL> CONNECT TO BankTranx user1;
dmSQL> CREATE DATABASE LINK LK1 CONNECT TO BankMIS IDENTIFIED BY user2;
```

```
dmSQL> CREATE VIEW v1 FOR BankTranx:Account;
dmSQL> CREATE VIEW v2 FOR LK1:user3.Personnel;
dmSQL> SELECT * FROM v1;
       // SELECT * FROM BankTranx:user1.Account; (BankTranx, user1)
dmSQL> SELECT * FROM v2;
       // SELECT * FROM BankMIS:user3.Personnel; (BankMIS, user2)
dmSQL> DISCONNECT;
dmSQL> CONNECT TO BankTranx user3;
dmSQL> CREATE DATABASE LINK LK1 CONNECT TO BankMIS IDENTIFIED BY user4;
dmSQL> SELECT * FROM v1;
       // SELECT * FROM BankTranx:user1.Account; (BankTranx, user1)
dmSQL> SELECT * FROM v2;
       // SELECT * FROM LK1:user3.Personnel; (BankMIS, user2)
```

The comments indicate the equivalent SQL expression, the database being connected to, and the account used to connect.

Closing Links

Once a user accesses a remote database with an SQL command, the Coordinator Database will build a remote connection to the Participant Database. The remote connection will remain open until all users disconnect from the Coordinate Database or until the link is closed with the CLOSE DATABASE LINK command. DBMaker provides up to 256 remote connections for each database; it is a good idea to close remote connections that are no longer in use to free the connections for other users.

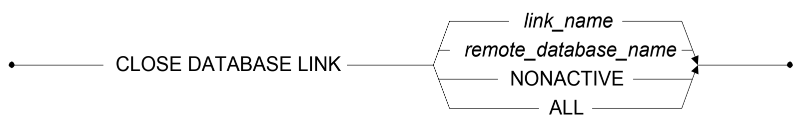

Figure 15-6 Syntax for closing a database link

➲ **Example 1**

To close a database link using the remote database name **BankMIS**:

```
dmSQL> close database link BankMIS;
```

➲ **Example 2**

To close a database link using the remote database link **BankLink1**:

```
dmSQL> close database link BankLink1;
```

When a user issues a CLOSE DATABASE LINK command, DBMaker will decrease the remote connection counter by one. When the counter reaches zero, the connection is fully closed and the occupied resources freed. Otherwise, the connection remains open.

➲ **Example 3**

To close all database links and free the connections and resources:

```
CLOSE DATABASE LINK ALL
```

➲ **Example 4**

To close all NONACTIVE remote connections no longer being used:

```
CLOSE DATABASE LINK NONACTIVE
```

Link System Catalog Tables

Two system catalog tables relate to database links: SYSDBLINK and SYSOPENLINK. SYSDBLINK records all database link names and their definitions, while SYSOPENLINK, records open connections between databases.

15.5 Distributed Transaction Control

DBMaker supports a distributed transaction mechanism that is transparent to users. Participant Databases do not commit only a part of a distributed transaction; DBMaker handles it.

➲ **Example**

The following shows how distributed transaction control works.

```
dmSQL> CONNECT TO BankTranx user1;                    // ABC Bank in Taipei
dmSQL> SET AUTOCOMMIT OFF;
dmSQL> UPDATE BankTranx:Customer SET money=money-1000 where id=123;
dmSQL> UPDATE BankTranx@"Bank_in_Seattle":Customer SET money=money+1000
```

```
    2> WHERE id=123;
dmSQL> COMMIT;
```

Since the Coordinator Database handles all database operations from the client application, the Coordinator Database knows the scope of the instructions via the Distributed System Catalog Manager. The Coordinator Database will handle transactions belonging to the local site in the same manner as a regular client/server transaction. Transactions belonging to remote sites need to reference the appropriate remote database. The Coordinator Database will exchange information with every Participant Database and coordinate the whole transaction until it is either rolled back or committed.

Two-Phase Commit

Database management systems need to maintain data integrity, and this requires that all transactions be *atomic*. All operations in the transaction must commit or roll back together. In the traditional client/server architecture, a journal is used to make sure that the changes are either rolled back or committed.

In the distributed database architecture, a two-phase commit protocol with presumed abort is used as the mechanism for controlling distributed transactions that span multiple database servers. A transaction that modifies data on two or more databases must complete the two-phase commit protocol before it is committed. The two-phase commit mechanism guarantees that all sites commit or roll back globally. It also protects data manipulation operations performed by remote synonyms, integrity constraints, and triggers. To commit a transaction, a user has to ensure every sub-transaction has finished; otherwise, the transaction will be aborted. For the same reason, if any sub-transaction cannot commit, the other sub-transactions must be aborted as well.

Distributed Transaction Recovery

DBMaker uses the two-phase commit protocol to inform all Participant Databases to commit a global transaction. Before entering the commit phase, the Coordinator Database will check the status of the Participant Databases to ensure there are no server or network problems. If the Coordinator Database finds a problem with any of the Participant Databases, it informs the other Participant Databases to roll back their part of the transaction, and returns an error indicating the global transaction has failed. If the two-phase commit protocol has finished the preparation phase, but there is a server or network problem with a Participant Database, the global transaction is regarded as a success. DBMaker records which Participant Database server cannot commit its part of the transaction in the SYSGLBTRANX system catalog table, and records which database contains a pending transaction in the SYSPENDTRANX system catalog table for the crashed database.

DBMaker also provides an automatic recovery mechanism to handle network or site failure during the execution of a distributed transaction. In the Coordinator Database, you start the global transaction recovery daemon (GTRECO). This daemon scans the SYSGLBTRANX system catalog table and periodically recovers any pending global transactions. Then it tries to connect to the crashed Participant Database and informs it to commit or roll back its part of the global transaction.

Heuristic End Global Transaction

After a network or site failure occurs during the two-phase commit, pending transactions continue to hold some resources such as locks or journal blocks. The pending transaction will occupy these resources until the problem is solved by the global recovery daemon (GTRECO). If the network or site failure cannot be solved immediately, then the held resources may block some of the users in the Participant site. To solve this problem, DBMaker supports heuristic end global transaction. A heuristic end transaction is an independent action taken by the database administrator to force a pending transaction in a Participant Database to commit or roll back. The database administrator can use JDBATool to solve this problem. Refer to the *JDBA Tool User's Guide* for more information.

➲ **To resolve a pending transaction manually, perform the following:**

1. In the Participant Database, browse the SYSPENDTRANX table to find out whether there are transactions that have been pending for a long time.
2. In the Coordinator Database, determine the commit status of the pending transaction from SYSGLBTRANX. Also determine the commit status of the pending transaction from the Participant Database. For example, if there are two

pending transactions "DB_1-3376aafd" and "DB_2-3376aafd:DB_3-3376ab0f#1", the administrator of the coordinator database should ask the administrator of DB_1 to determine the status of "DB_1-3376aafd" and ask the administrator of DB_3 to determine the status of "DB_2-3376aafd:DB_3-3376ab0f#1".

3. In the Coordinator Database, when the administrator receives the transaction status query he can examine SYSGLBTRANX to determine the transaction status. If the STATE is 2 (COMMIT) or 3 (PENDCOM), reply “commit”. If the STATE is 4 (PENDABO), reply ‘abort’. But if the STATE is 1 (PREPARE), this transaction branch is pending in this site too, and ask the administrator of the parent site to determine its status.

4. In the Participant Database, the administrator uses JServer Manager to perform a heuristic commit or abort the pending transaction based on the reply.

If the administrator initiates a heuristic end transaction on a pending transaction that is different from action taken in the Coordinator Database, the distributed data will be inconsistent.

16 Data Replication

Data Replication, in the broadest sense of the term, refers to the process of representing objects in more than one database.

Only a few years ago, corporate data resided in a central location. Remote departments accessed the information they needed by establishing direct connections to the central sites, or by requesting printed reports from central MIS. The connections however were expensive, unreliable, and limited in number, while the reports were inflexible and untimely.

Open systems brought inexpensive and powerful computing resources to all corners of an enterprise. The ability to share corporate information effectively using these new resources became an important competitive advantage for organizations. The question enterprises face today is not "Why distribute and share corporate data?" but rather "How can one distribute information effectively?" Replication is quickly becoming the architecture of choice for a majority of distributed corporate applications.

16.1 Table Replication

What is Table Replication?

Table replication creates a full or partial copy of a table to a destination location. This allows users in remote locations to work with a local copy of data. The local copy remains synchronized with the databases in other locations. This way each database can service data requests immediately and more efficiently, without having to go to another machine over a slower wide area network connection. This kind of request happens frequently between the headquarters and area companies or branch offices.

For example, after creating a replication from table A on Taipei's database server to table B on Tokyo's database server, modifications made to table A will replicated to table B. Clients in Tokyo can then access Tokyo's database rather than connecting to Taipei's database to acquire the same data.

Destination tables may also reside on databases on the same server. This situation may occur if data must be shared between two databases designed for different functions, or in the case of sharing data between databases of different types (for example, Oracle, Sybase, etc.). Tables that are receiving data from another database through replication will be referred to as destination tables rather than remote tables, even though the destination tables may reside on a remote database.

Differences Between Database and Table Replication

The major difference between database and table replication is that the replicated data object is different: one is the whole database; the other one is a table. Users can choose one of them depending on their demands. If you choose database replication, since the unit to replicate is the whole database, the target (or slave) database is read-only.

Two Types of Table Replication

There are two types of table replication. One is synchronous. 'Synchronous' means the modification to the destination site shows up immediately; the destination table is modified at the same time as the local table. DBMaker uses a 'two phase commit' and triggers to perform synchronous table replication. Thus, after establishing a replication, any update on the source table will become a DDB (distributed database) action. This will affect the local database's behavior. If the destination database server is unreachable, updates on the local database will fail.

The other table replication type is asynchronous. Modifications to the destination site are delayed. The delay between source and destination database depends on a user-defined schedule. Asynchronous table replication stores changes to the local table and modifies the destination table based on a schedule. In this type of replication, two databases of a replication pair are independent and can work as normal even if the network is not available.

Term Definitions

SOURCE TABLE

The table on the source database that the data is replicated from.

DESTINATION TABLE

The table on the destination database that the data is replicated to.

PUBLICATION

A data set on the source table that is available for replication.

SUBSCRIPTION

The data set on the destination table to receive a publication.

FRAGMENT

Also called a horizontal partition, a fragment is the replication of a given range of data.

PROJECTION

The selected columns from a base table chosen for replication.

REPLICATION DOMAIN

A replication fragment (horizontal partition) and projection (vertical partition) are called a replication domain. It is the range of a table's data to be replicated.

There is a problem when replicating a domain change. It may occur when replicating an UPDATE statement.

➲ **Example**

```
CREATE REPLICATION rp2 WITH PRIMARY AS t1 WHERE c2 > 0 REPLICATE TO db2:t1;
CREATE REPLICATION rp3 WITH PRIMARY AS t1 WHERE c2 < 0 REPLICATE TO db2:t1;
```

```
UPDATE t1 SET c2=-7 WHERE c2=7;
```

The **rp2** replication domain is c2>0, and the **rp3** replication domain is c2<0. When replicating, it is not only replicating the UPDATE statement. The updated tuples replication domain is changed from **rp2** to **rp3**, subsequently the replication performs a DELETE statement for **rp2** (DELETE c2=-7) and performs an INSERT statement for **rp3** (INSERT c2=7).

DATA INITIALIZATION

When creating replications, users can specify how to initialize the data on both sides automatically. After creating a table replication, any modification (insertion, deletion, update), to a source table will affect the destination tables.

DBMaker provides 4 options:

- **Clear data** — DBMaker will delete all data from the destination table while performing table replication.
- **Flush data** — DBMaker will insert all data tables that satisfy the fragment for the source table into the destination tables.
- **Clear and flush data** — DBMaker will 'clear data' and then 'flush data'.
- **Do nothing** — Keeps the destination tables.

Creating Table Replication

The following syntax diagram describes both asynchronous and synchronous table replication.

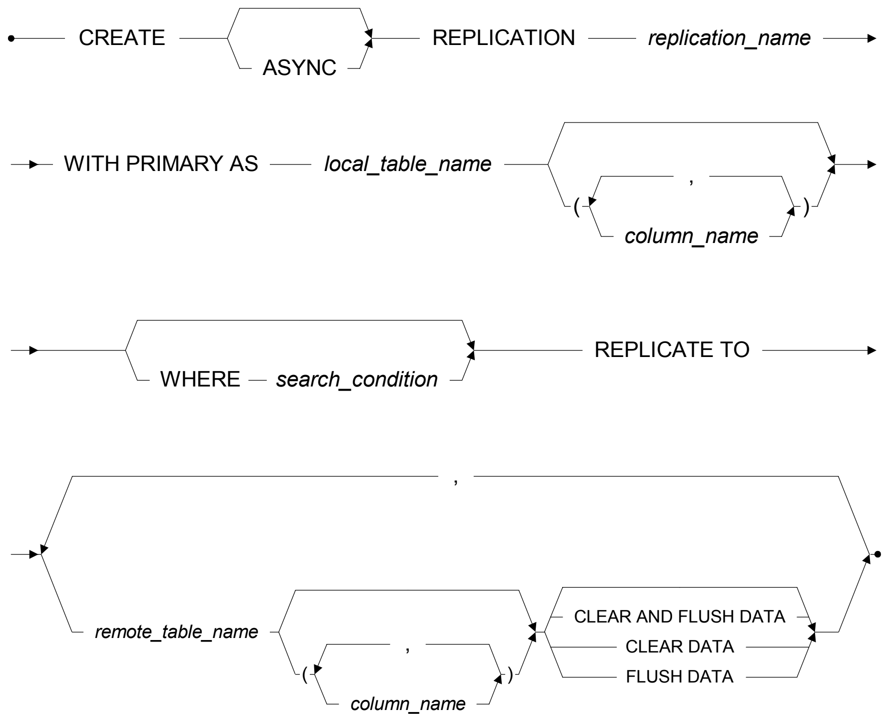

Figure 16-1 Syntax for the CREATE REPLICATION Statement

➲ **Example 1**

Suppose there is a table TB1 in database DB30A, a table TB2 in DB30B. We want to replicate TB1's data to TB2 and do not want to modify TB2's current data.

```
dmSQL>  CREATE REPLICATION r1 WITH PRIMARY AS TB1
            REPLICATE TO DB30B:TB2;
```

In the above example, we use a database session name to specify the destination database. Alternatively, a database link name could be used.

Table Replication Rules

- Schemas of source and destination tables must exist. This means DBMaker does not create tables when performing table replication.
- Replication names for a table must be unique.
- The subscriber name for a replication must be unique using the <*link_name*|*database_session_name*>+ <*table_owner_name*> + <*table_name*> syntax.
- Projected columns for every table must contain primary key columns.
- Primary key columns must be included with fragment columns.
- Only the table owner of the source table or a user with DBA or higher authority has the privilege to create, drop, or alter replications.
- If no column-identifier exists with the destination table, the destination column names must be the same as the base table names.
- Number of primary key columns must be equal in the base table and destination tables.

➲ **Example 1**

The following statement will create a publication that will replicate table **t1** from the local database to table **usr1.t2** on **db1**. Without specifying column names, all columns in table **t1** must exist in **t2** and column types must be compatible. If **t1** contained 3 columns **c1**, **c2**, and **c3**, it would replicate **c1**, **c2**, and **c3** from table **t1** to columns **c1**, **c2**, and **c3** on table **usr1.t2**.

```
dmSQL> CREATE REPLICATION r1 with
       PRIMARY AS t1
       REPLICATE TO db1:usr1.t2;
```

➲ **Example 2**

The following statement will create a publication used to replicate tuples where **c1** > 100 with columns (**c1**, **c2**) to (**column1**, **column2**) in **t2** on **db1**, and (**c1**, **c2**) in **t3** on

db2. After issuing this command, all data where **c1** > 100 on **t1** will be replicated to **t3** on **db2** and **column1** and **column2** of **t2** on **db1**:

```
dmSQL> CREATE REPLICATION r2 with
                PRIMARY AS t1 (c1,c2) where c1 > 100 ,
        REPLICATE TO db1:t2 (column1, column2),
                    db2:t3 flush data;
```

➲ **Example 3**

The first function of this command is to delete all data from **t2** on **db1** and then to create a publication used to replicate the records where **c1** > 100 with only column (**c1**, **c2**) to (**column1**, **column2**) in **t2** on **db1**:

```
dmSQL> CREATE REPLICATION r2 with
                PRIMARY AS t1 (c1,c2) where c1 > 100 ,
        REPLICATE TO
                db1:t2 (column1, column2) clear data;
```

Drop Replication

This command will drop a replication from a source table.

•— DROP REPLICATION — *replication_name* — FROM — *table_name* —•

Figure 16-2 Syntax for the DROP REPLICATION Statement

➲ **Example**

To drop the replication **r1** from the **t1** table:

```
dmSQL> DROP REPLICATION r1 FROM t1;
```

Alter Replication

Users may add destination tables to or drop destination tables from an existing table replication. The following syntax diagrams and examples demonstrate how.

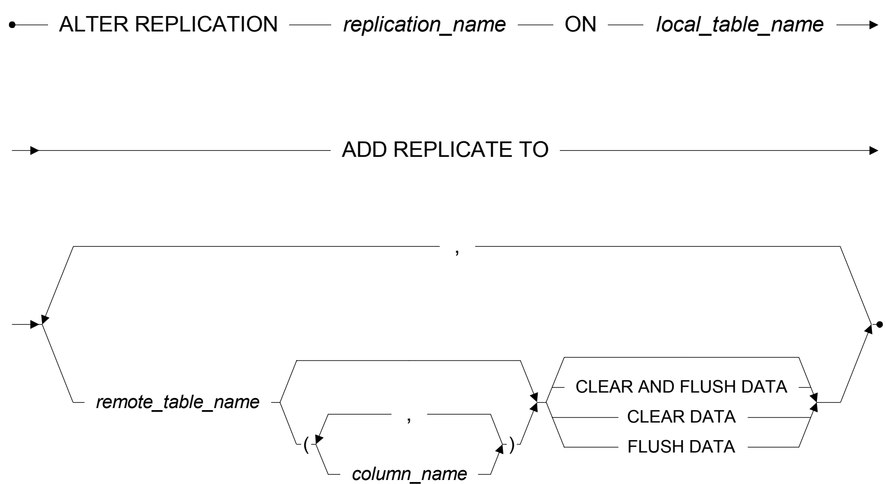

Figure 16-3 Syntax for the ALTER REPLICATION … ADD REPLICATE TO Statement

➲ **Example 1**

The first command creates a replication to replicate table **t1** on the local database to table **t** on database **dbX**. The second command adds 2 subscribers, **t3** on **db1** and **tableA** on **db4**, to the replication of **r3** in table **t1**.

```
dmSQL> CREATE REPLICATION r3 WITH PRIMARY AS t1
       REPLICATE TO dbX:t;
dmSQL> ALTER REPLICATION r3 ON t1
       ADD REPLICATE TO db1:t3,
                   db4:tableA (col1, col2) clear data;
```

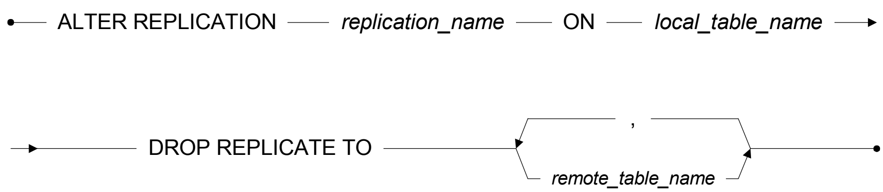

Figure 16-4 Syntax for the ALTER REPLICATION … DROP REPLICATE TO Statement

➲ **Example 2**

To drop subscriber **t3** on **db1** from the replication of **r3** for table **t1**:

```
dmSQL> ALTER REPLICATION r3 ON t1
       DROP REPLICATE TO db1:t3;
```

16.2 Synchronous Table Replication

Two-phase commit allows the synchronization of distributed data. A transaction will only be accepted if all interconnected distributed sites agree. A 'handshake' mechanism across the network allows distributed sites to coordinate their acceptance for each transaction. Therefore, using synchronous table replication guarantees that data will be synchronous whenever an update occurs.

Synchronous Table Replication Setup

DBMaker uses two-phase commits to perform synchronous table replication. Source and destination databases must be in distributed database (DDB) mode. Therefore, the first step is to start the databases in DDB mode (**DD_DDBMd** =1) and add database sessions to the **dmconfig.ini** file. Refer to Chapter 15, *Distributed Databases* for more information.

After creating a table replication, any modifications (insertion, deletion, updates) to source a table will affect the destination tables.

16.3 Asynchronous Table Replication

Synchronous table replication modifies the destination table at the same time it modifies the local table, while asynchronous table replication stores changes to the local tables and modifies the destination table based on a schedule.

DBMaker uses a file known as a replication log to store changes to local tables. Modifications to local tables are stored in replication logs, and are replicated to the destination table according to a predefined schedule. Using replication logs enables DBMaker to treat the local transaction and the destination transaction independently, allowing you to update local tables normally even if the remote connection is not available. This allows asynchronous table replications to tolerate network and destination database failures, since DBMaker will keep trying until failures are corrected.

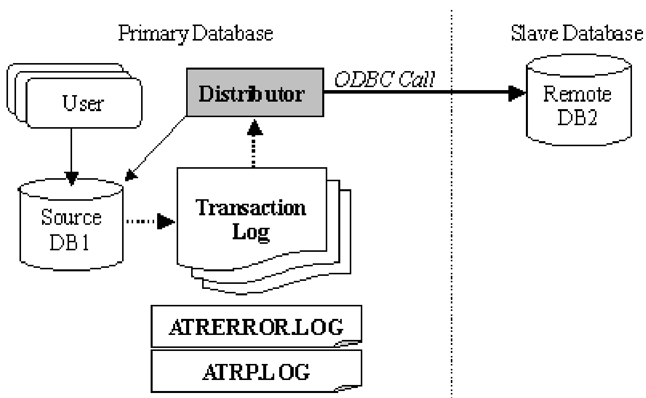

Figure 16-5: Architecture of asynchronous table replication

Asynchronous table replication uses a replication log system and the distributor server to handle data replication. Replication logs are not DBMaker journal files. The level of the replication log is higher than the journal, and it is only for replicating tables. The content of a journal is physical data modification, but the replication log consists of commands that are applied to the destination tables.

When the source database is running, DBMaker logs the modification of source tables into the replication log files. When the distributor server activates, it will redo all changes made to the source tables to the destination tables according to the replication log.

Normally the distributor server uses ODBC function calls to communicate with the destination database servers, so it is possible to replicate tables to heterogeneous database servers such as Oracle, SQL Server, Informix, etc. Heterogeneous table replication will be covered later in this chapter. Express asynchronous table replication is another type of asynchronous table replication that does not use ODBC function calls. It will also be covered later in this chapter.

➲ **There are three main steps to building an asynchronous table replication:**

1. Enable asynchronous table replication.
2. Create a schedule to the destination database.
3. Based on the schedule, create the asynchronous table replication.

Enabling Asynchronous Table Replication

The distributor server resides in the source database. The distributor connects to the destination databases periodically and performs the table replication.

The **DB_AtrMd** keyword in the **dmconfig.ini** file in the source database specifies whether to start the distributor server for a database. If you do not start the distributor server, the database cannot be the source for asynchronous table replications.

The **RP_LgDir** keyword in the **dmconfig.ini** of the source database specifies the directory where DBMaker will place replication log files for asynchronous table replication. The replication log files are binary and users should not manually remove them. The default directory of **RP_LgDir** is the subdirectory named **/TRPLOG** in the database home directory.

When creating table replications with schema checking, the distributed database mode in the source and the destination database must be enabled (**DD_DDBMd** = 1). If the schedule is set to NO CHECK, DBMaker will not check the schema, and the distributed database mode can be set to OFF (**DD_DDBMd** = 0). See the following section for more information about creating replication schedules.

➲ **Example 1**

To replicate a table from the **SRCDB** database to the destination **DESTDB** database, add the following lines to the **dmconfig.ini** file on the source database server:

```
[SRCDB]
DB_DBDIR = /disk1/DBMaker/src
DB_USRBB = /disk1/DBMaker/src/SRCDB.BB 2
DB_USRDB = /disk1/DBMaker/src/SRCDB.DB 150
RP_LGDIR = /disk1/DBMaker/src/trplog
DB_ATRMD = 1
DD_DDBMD = 1
DB_SVADR = srcpc
DB_PTNUM = 22222

[DESTDB]
DB_SVADR = destpc
DB_PTNUM = 33333
```

The **dmconfig.ini** file on the destination database:

```
[SRCDB]
DB_SVADR = srcpc
DB_PTNUM = 22222

[DESTDB]
DB_DBDIR = /disk3/DBMaker/dest
DB_USRBB = /disk3/DBMaker/dest/DESTDB.BB 2
DB_USRDB = /disk3/DBMaker/dest/DESTDB.DB 150
DD_DDBMD = 1
DB_SVADR = destpc
DB_PTNUM = 33333
```

Due to the way distributor server makes use of the ODBC driver manager to perform asynchronous table replication, the ODBC data source name (DSN) of the destination database must be set if the source database is running under the Microsoft Windows environment.

Schedule (Creating and Dropping)

Before creating asynchronous replication to one or more destination tables, a user with DBA privilege or higher needs to define a *schedule*. A schedule defines the starting time, the period, and the account and password used to connect to the destination database. The user can create several schedules for different destination databases on the same source database, but cannot build more than one schedule to the same destination database.

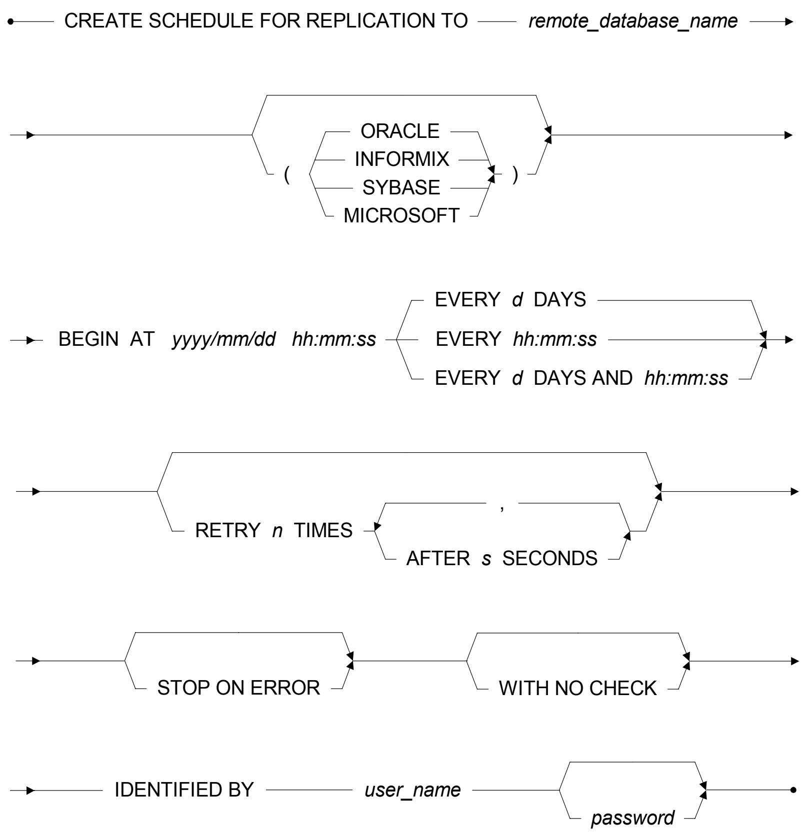

Figure 16-6 Syntax for the CREATE SCHEDULE Statement

➲ **Example 1**

To create a schedule for the **DESTDB** database:

```
dmSQL > CREATE SCHEDULE FOR REPLICATION TO destdb
         BEGIN AT 2000/1/1 00:00:00
         EVERY 12:00:00
         IDENTIFIED BY User Password;
```

From January 1, 2000, the distributor server will activate every 12 hours to perform an asynchronous replication. The *ATRP.LOG* distributor message log in the database home directory will record the starting time and status of the distributor server.

The IDENTIFIED BY keyword specifies the destination account used by the distributor server to connect to the destination database and execute the table replication. The account must have privilege to insert, delete, and update on the destination tables.

If the schedule is not necessary and there is no replication associated with it, users having the DBA privilege on the source database can drop it.

•—— DROP SCHEDULE FOR REPLICATION TO —— *remote_database_name* ——•

Figure 16-7 Syntax for the DROP SCHEDULE Statement

➲ **Example 2**

To drop a schedule:

```
dmSQL> DROP SCHEDULE FOR REPLICATION TO destdb;
```

Creating Asynchronous Table Replication

All asynchronous table replications depending on the same schedule will be done at the same time. The asynchronous table replication mechanism supports all data types, including LONG VARCHAR, LONG VARBINARY, and FILE types.

The action to create an asynchronous table replication is similar to creating a synchronous table replication. Add a keyword (ASYNC, CREATE ASYNC, REPLICATION command), to create table replication based on one schedule to a destination database.

➲ **Example 1**

To create a replication named rp1 for the **SRCDB** database, based on the schedule to the **DESTDB** destination database:

```
dmSQL> CREATE ASYNC REPLICATION rp1
       WITH PRIMARY AS t1
       REPLICATE TO destdb:t2;
       CLEAR AND FLUSH DATA;
```

Users can specify CLEAR DATA, FLUSH DATA, or CLEAR AND FLUSH DATA to perform data initialization for a destination site. When creating replication, DBMaker may use a database link to connect to the destination database for checking and initialization. The action is performed through the current user account or the destination account using the DESTDB destination database or database link name.

After the asynchronous table replication has been created, the job of replicating is moved to the distributor server. The account used to connect to the destination database will be changed to the one specified by the IDENTIFIED option defined in the CREATE SCHEDULE statement.

Any transactions resulting from the asynchronous table replication will be recorded in the distributor message log **ATRP.LOG** under the source database home directory. The distributor message log is a pure text file, and records the start-up and actions of the distributor server.

➲ **Example 2**

Typical content of an ATRP.LOG file:

```
2000/02/09 10:02:30 : start up
2000/02/09 10:02:33 : replicate transactions before 2000/02/09
                      10:02:29 (log:1.856152) to DESTDB
```

NO CASCADE OPTION

The NO CASCADE keywords are optional. It only functions when the replication type is asynchronous. The keyword specifies cascade replication. Commands flow in most organizations from the highest level to lower levels; for example, replicating data from A to B, and then B to C. This is a typical kind of cascade replication. A typical no-cascade model replicates data to B and B replicates data to A. If your data model works like this, you can turn on the NO CASCADE option. The default is setting is CASCADE.

The NO CASCADE option causes table replication to occur only across one site.

➲ **Example**

Replication **Rp1** is from **DB1:t1** to **DB2:t2**, and **Rp2** is from **DB2:t2** to **DB3:t3**. If **Rp1** has the NO CASCADE option set, the changes for **DB1:t1** will be replicated to **DB2:t2**, but **DB2** will stop replicating the same changes to **DB3:t3**. If **Rp1** has the CASCADE option, the changes of **DB1:t1** will be replicated to **DB2:t2**, and then from **DB2:t2** to **DB3:t3**.

Error Handling

In the process of data replication, the distributor may face five kinds of errors: warning, connection, data, statement, and transaction.

Warning

For example, if a CHAR(10)data type is replicated to a CHAR(5) column type, there will be a data truncation warning. The distributor server will ignore these kinds of errors.

Connection Errors

If the distributor server fails to connect to the destination database server, it will abandon the schedule and wait until the next time. All jobs will be kept until then.

Data Errors

Because asynchronous table replication is loosely coupled, it is possible that someone else will update the destination data first. For example, the distributor server wants to insert a record to a destination database, but it already exists. Another situation could be that the distributor server wants to delete a record, but it does not exist. Users can use the STOP ON ERROR option to make the distributor server stop when it incurs such errors. The default behavior of the distributor is to ignore these kinds of errors.

NOTE *The data error mentioned above includes an integrity violation error and affected row error. The integrity violation error calls function SQLError() and returns the '23000'error state. The affected row error calls function SQLRowCount(), the result is not one. For more information on ODBC functions, refer to the "ODBC Programmer's Guide".*

STATEMENT ERRORS

When the distributor server faces lock time-out errors when executing a statement, it needs to wait and retry using the RETRY <*n*> TIMES option. The AFTER <*s*> SECONDS option specifies how long to wait before the next try.

TRANSACTION ERRORS

For example, a dead lock causes transactions belonging to transaction errors to be rolled back. If the distributor server faces errors that rollback transactions, it will retry once for each whole transaction. If the outcome still fails, the distributor will leave these actions for the next schedule.

Users can check a text file named **ATRERROR.LOG** for error records that occurred during replication. This file is located in the home directory of the database.

Schedule (Suspending and Resuming)

If we replicate data to a database located in Tokyo, and we know that the database will be shut down on traditional holidays, we can suspend the schedule until the database is ready.

Users with the DBA privilege can suspend and resume schedules. After a schedule is suspended, the distributor server will stop trying to connect for replications.

➲ **Example 1**

To suspend the schedule to the **DESTDB** destination database:

```
dmSQL> SUSPEND SCHEDULE FOR REPLICATION TO destdb;
```

➲ **Example 2**

To restart the schedule for the **DESTDB** destination database:

```
dmSQL> RESUME SCHEDULE FOR REPLICATION TO destdb;
```

Synchronizing a Replication

Users will need synchronized data sometimes; to achieve this DBMaker provides synchronized schedules. The synchronized schedule forces the distributor server to perform local changes to the specified database immediately. Users do not need to wait until the schedule activates the distributor server.

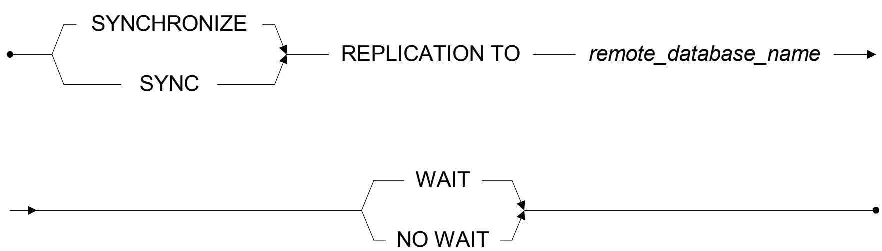

Figure 16-8 Syntax for the SYNCHRONIZE REPLICATION Statement

➲ **Example 1**

To synchronize schedule to the **DESTDB** destination database:

```
dmSQL> SYNC REPLICATION TO destdb WAIT;
```

The default WAIT option causes the distributor server to wait until all changes have been made. The command returns only after the completion of replication. The NO WAIT option instructs the distributor server to perform its job immediately, and the SYNC command will return immediately.

➲ **Example 2**

Using SYNC REPLICATION TO with NO WAIT:

```
dmSQL> SYNC REPLICATION TO destdb NO WAIT;
```

Altering Schedule

After creating a schedule, users having the DBA privilege can alter the attributes of a schedule, including the distributor's activation interval, the account used to connect to the destination database, the RETRY option and STOP/IGNORE ON ERROR option.

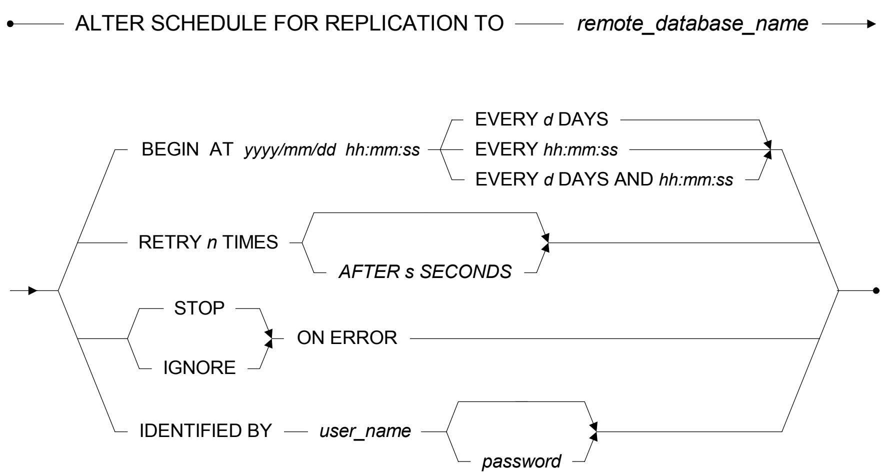

Figure 16-9 Syntax for the ALTER SCHEDULE Statement

- **Example 1**

To alter the schedule for replications to the **DESTDB** destination database by adding the IGNORE ON ERROR option:

```
dmSQL> ALTER SCHEDULE FOR REPLICATION TO destdb IGNORE ON ERROR;
```

- **Example 2**

```
dmSQL> ALTER SCHEDULE FOR REPLICATION TO destdb
       IDENTIFIED BY User2 Password2;
dmSQL> ALTER SCHEDULE FOR REPLICATION TO destdb STOP ON ERROR;
dmSQL> ALTER SCHEDULE FOR REPLICATION TO destdb RETRY 5 TIMES AFTER 3
       SECONDS;
```

Heterogeneous Asynchronous Table Replication

DBMaker not only allows asynchronous table replication to other DBMaker databases, but also to Oracle, Informix, Sybase, and Microsoft SQL Server databases. This type of replication allows DBMaker to coexist with other databases in a heterogeneous environment, and is known as *heterogeneous table replication*.

DBMaker needs to preprocess the replicated data before sending it to a third-party destination database; users must specify the type of DBMS they are replicating to when creating a schedule in a heterogeneous environment using the ORACLE, INFORMIX, SYBASE, and MICROSOFT keywords.

Due to the way DBMaker makes use of the ODBC Driver Manager to perform asynchronous table replication, the DBMaker server must be located on a computer running Windows NT or Windows 2000, and the definition of the destination database name cannot include a link name. The third-party destination databases may be located on either Windows, UNIX, or Linux platforms.

When creating a schedule for heterogeneous table replication, use the WITH NO CHECK keywords to prevent DBMaker from performing schema checking. The user creating the replication must take responsibility for schema checking, and ensure that columns and data types in the destination table are compatible with the columns and data types in the local table.

➲ **Example 1**

To connect to an Oracle database with an ODBC data source name of **orcdb**, user **orcuser** with password **mypassword** enters:

```
dmSQL> CREATE SCHEDULE FOR REPLICATION TO orcdb (ORACLE)
       BEGIN AT 2000/01/01 00:00:00 EVERY 2 DAYS
       WITH NO CHECK
       IDENTIFIED BY orcuser mypassword;
```

The CLEAR DATA, FLUSH DATA, or CLEAR AND FLUSH DATA keywords cannot be used when creating a heterogeneous table replication. Data in the third-party destination database must be manually deleted or inserted to put the table in its initial state before the replication begins. Heterogeneous replication schedules use the same syntax as homogeneous schedules.

➲ **Example 2**

The following shows the heterogeneous replication of table **t1**:

```
dmSQL> CREATE ASYNC REPLICATION rp1
       WITH PRIMARY AS t1
       REPLICATE TO orcdb:orcuser.t1;
```

Express Asynchronous Table Replication

Asynchronous table replication uses ODBC function calls to communicate with destination databases, which might cause poor performance in a WAN environment. To achieve better performance on a WAN, DBMaker provides another mechanism named express asynchronous table replication. DBMaker packs commands into a package to travel the network.

Since other database management systems do not support this protocol, express asynchronous table replications cannot work with on heterogeneous replications. It also does not support the STOP ON ERROR option when creating express schedules.

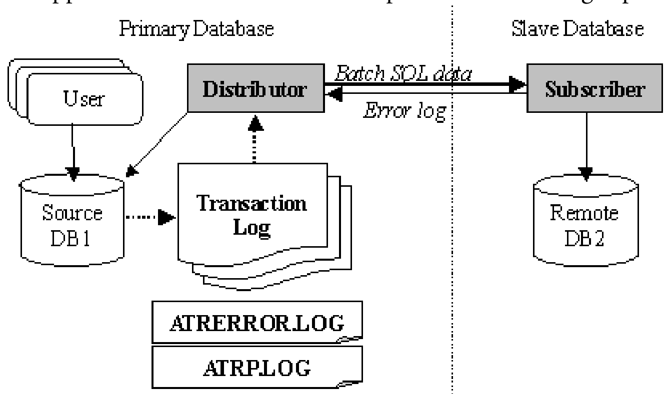

Figure 16-10: Architecture of express asynchronous table replication

There is a distributor server on the source database and a subscriber daemon on the destination database. They co-operate to do the work for express asynchronous table replication. The distributor does not directly apply the changes from the source table to the destination table via ODBC calls. Instead, it only packages the SQL commands and related data applied on the source table, and sends the packet out to the subscriber daemon on the destination database. After getting the packet on the destination database, the subscriber daemon applies the commands to the destination tables.

Express Replication Setup

➲ **To build an express replication:**

1. Enable the distributor server on the source database and subscriber daemon on the destination database(s).
2. Create an express schedule to the destination database(s).
3. Based on the schedule, create the asynchronous table replication(s).

ENABLING SUBSCRIBER DAEMON

In order to start the subscriber daemon, the **DB_EtrPt** keyword in the **dmconfig.ini** file for the destination (subscriber) database should be set. It specifies the port number of the communicating channel between the distributor server and the subscriber daemon.

➲ **Example 1**

To replicate a table from the **SRCDB** source database to the **DESTDB** destination database by express replication, the subscriber daemon should be started on the destination database.

A sample **dmconfig.ini** file in a destination database follows:

```
[SRCDB]
DB_SVADR = srcpc    ; tell the target database where the source
DB_PTNUM = 22222    ; database is
[DESTDB]
DB_DBDIR = /disk3/DBMaker/dest
DB_USRBB = /disk3/DBMaker/dest/DESTDB.BB 2
DB_USRDB = /disk3/DBMaker/dest/DESTDB.DB 150
DD_DDBMD = 1
DB_SVADR = destpc
DB_PTNUM = 33333
DB_ETRPT = 44444    ; port number used by Subscriber Daemon
```

In the source database, the **DB_AtrMd** keyword is used to start the distributor server. It is not necessary to let the distributor know which port number the subscriber daemon uses for the destination database.

A sample **dmconfig.ini** file from the source database follows:

```
[SRCDB]
DB_DBDIR = /disk1/DBMaker/src
DB_USRBB = /disk1/DBMaker/src/SRCDB.BB 2
DB_USRDB = /disk1/DBMaker/src/SRCDB.DB 150
RP_LGDIR = /disk1/DBMaker/src/trplog
DB_ATRMD = 1
DD_DDBMD = 1
DB_SVADR = srcpc
DB_PTNUM = 22222
[DESTDB]
DB_SVADR = destpc
DB_PTNUM = 33333
```

Schedule for Express Asynchronous Table Replication

Express asynchronous table replication uses the EXPRESS option specified in the CREATE SCHEDULE command.

➲ **Example**

To build a schedule for express asynchronous table replication:

```
dmSQL> CREATE SCHEDULE FOR EXPRESS REPLICATION TO destdb
       BEGIN AT 2000/1/1 00:00:00
       EVERY 12:00:00
       IDENTIFIED BY User Password;
```

Creating Express Asynchronous Table Replication

The steps are the same as for creating asynchronous table replications. Refer to the section *Creating Asynchronous Table Replication*, or the *SQL Command and Function Reference* for more information.

16.4 Database Replication

Most enterprises or companies had to put all data in one file server or database in their headquarters, and all terminals needed to connect to the server directly. Under this architecture, the application system might work smoothly if all terminals are in the same building or area. However, the performance was much slower if the remote branch wanted to access the database, because of transmitting speed and network bandwidth.

To share data and increase access speed, DBMaker provides database replication. The database replication will replicate the primary database to the slave database during a set time frame. In other words, when there are changes happening to the primary database like newly added, modified, or deleted data, DBMaker will replicate the changed data into the slave database automatically. There are three advantages to using database replication. First, the performance of application systems will increase since it can access data from the local side directly. Second, the destination database will have the same data modification as the local one; thus, the goal of data sharing can be achieved efficiently. Third, if it fails to connect to the local database for an application, it still could connect to the remote database and continue to work. Although there are advantages to using data replication, more storage for data and more processes to handle the replication are required.

Database Replication Basics

In this section, we will use Figure 16-11 to explain the flow of database replication. Database replication will work with the journal backup server, if you are not familiar with database backup and restoration, please read Chapter 14, "Database Recovery, Backup, and Restoration".

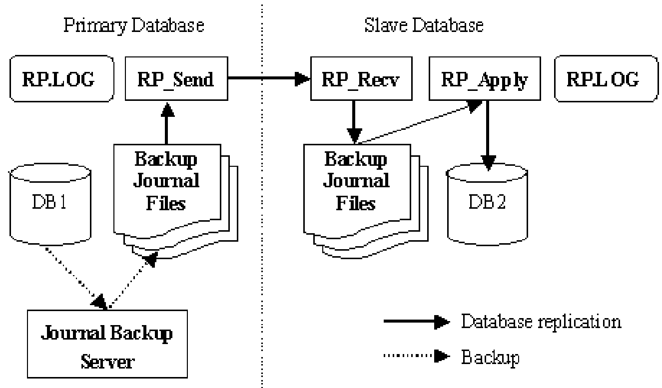

Figure 16-11: Flow of database replication

Database replication is accomplished with 4 servers: journal backup server, RP_Send server, RP_Recv server, and RP_Apply server. The journal backup server and the RP_Send server are started from the primary database. The RP_Recv server and the RP_Apply server are started from the slave database.

The initial step of database replication is to make the slave database the same as the primary. All subsequent changes, data insertion, deletion, creating schema, etc., have to be made in the primary database first. Then the journal backup server running on the primary database regularly writes changes to the backup journal.

The RP_Send server periodically sends the backup journal of the primary database to the slave- RP_Recv server. The RP_Recv gives notice to the RP_Apply to apply the received backup journal to the slave database.

The slave database is read-only for all users, only RP_Apply can update the slave database.

Database Replication Setup

➲ **To manually set up database replication:**

1. Duplicate the primary database to the slave database.
2. Set up the journal backup server and RP_Send server on the primary database.

3. Set up the RP_Recv server and RP_Apply server on the slave database.
4. Start the primary and slave databases.
5. Verify the replication log (RP.LOG) and error log (ERROR.LOG).

FULL BACKUP

As mentioned earlier, the first step of database replication is to make the slave database identical to the primary. Please note the byte ordering of primary and slave database machines must be the same. For example, if the primary database is running on a x86 machine, the slave database machine must be byte ordering-compatible with x86; Sun Sparc's byte ordering is different from x86's, so the slave database cannot run on a Sparc machine and vice versa.

➲ **To duplicate the primary database:**

1. Shut down the primary database.
2. Make a copy of the primary database's data files, BLOB files, and journal files.
3. Move the copy of files to the machine that will run the slave database.
4. Modify the slave database's **dmconfig.ini** configuration file to reflect the corresponding file directory changes.

Step 1 and step 2 combined make an offline full backup. Refer to Section *Offline Full Backups* for more detailed information.

Here we will illustrate how to perform an offline full backup manually. In this section, all examples assume an x86 processor and the primary database running on Linux or FreeBSD.

➲ **Example 1**

To copy the primary database's files, query the SYSFILE system table to get the logical names:

```
dmSQL> SELECT * from SYSFILE;
```

An excerpt from the **dmconfig.ini** file to view the physical file names:

```
[MYDB]     ;; Primary Database Configuration, CPU:x86, OS: FreeBSD
DB_DBDIR = /home/dbmaker/mydb
```

```
DB_USRBB = /home/dbmaker/mydb/MYDB.BB 2
DB_USRDB = /home/dbmaker/mydb/MYDB.DB 150
FILE1 = /home/dbmaker/mydb/data/FILE1.DB 50
FILE2 = /home/dbmaker/mydb/data/FILE2.DB 50
DB_JNFIL = JN1.JNL JN2.JNL
...
```

After shutting down the primary database, copy the primary database files to the machine to run the slave database. Here, the slave machine is an x86 running Windows NT. In the above example, the files to copy include system files **MYDB.SDB** and **MYDB.SBB**, default user files **MYDB.DB** and **MYDB.BB**, journal files **JN1.JNL** and **JN2.JNL**, and user-defined files **FILE1.DB**, **FILE2.DB**, etc.

Next, copy the primary database section in the **dmconfig.ini** configuration file into the **dmconfig.ini** on the computer serving the slave database. Then modify the slave database **dmconfig.ini** to ensure physical file names logically match, since the actual paths and path name conventions may be different for different machines.

An excerpt from the **dmconfig.ini** of the slave database may look like this:

```
[MYDB]     ;; Slave Database Configuration, CPU：x86, OS: MS Windows NT
DB_DBDIR = d:\dbmaker\db\mydb
DB_USRBB = d:\dbmaker\db\MYDB.BB 2
DB_USRDB = d:\dbmaker\db\MYDB.DB 150
FILE1    = d:\dbmaker\db\mydb\FILE1.DB 50
FILE2    = d:\dbmaker\db\mydb\FILE2.DB 50
DB_JNFIL = JN1.JNL JN2.JNL
...
```

DBMaker supports up to 256 slave databases; if there is more than one slave database, the database administrator will have to repeat the above steps for each slave database.

Set Up Primary Database's Journal Backup Server

All changes to a database are logged to the journal files. DBMaker therefore uses backup journal files as the source data for replication. After the journal backup server performs an incremental backup, the RP_Send server sends backup journal files to the slave database. All the transactions logged on the journal files are then committed to ensure that all the changes to the primary database will also take place on the slave side.

➲ **Example**

Only the primary database needs to start the journal backup server. The **dmconfig.ini** file settings on the primary database server specify the backup directory and schedule:

```
DB_BMODE = 1                       ; Database start up in BACKUP-DATA mode
DB_BKSVR = 1                       ; Start up journal backup server
DB_BKDIR = /home/dbmaker/mydb/bkdir ; Directory of backup journal files
DB_BKTIM = 00/01/01 00:00:00       ; The initial backup time
DB_BKITV = 0-12:00:00              ; Perform journal backup evey 12 hours
```

DB_BMode can be BACKUP-DATA (1) or BACKUP-DATA-AND-BLOB (2) mode. However, even if **DB_BMode** is BACKUP-DATA-AND-BLOB mode, the BACKUP BLOB ON option still must be set for all tablespaces to be replicated, otherwise BLOB data will not be replicated to the slave database.

DATA TRANSMITTED AND RECEIVED

In the process of database replication, there are three resident servers RP_SEND, RP_RECV and RP_APPLY involved (as illustrated in Figure 16-11). RP_SEND is located on the primary database and is responsible for sending backup journal files to the slave side. The RP_RECV and RP_APPLY are on the slave database side. RP_RECV receives backup journal files from the primary database, and RP_APPLY executes the changes according to the journal files. RP_SEND is started and shut down automatically whenever the primary database is started or shut down. RP_RECV and RP_APPLY also start and end simultaneously with the slave database.

After the journal backup server completes an incremental backup, RP_SEND will send all the backup files that have not been sent to the slave database yet from the backup directory (**DB_BkDir**) to the slave database. Meanwhile, RP_RECV on the slave database will receive all the backup files transferred from the primary database, and put these files under the backup directory (**DB_BkDir**) on the slave database. After the files have been received, RP_APPLY will execute the changes recorded in the backup journal files to the slave database, and the flow of database replication will be complete.

SET UP RP_SEND SERVER ON THE PRIMARY SIDE

RP_SEND and RP_RECV are responsible for backup journal file transferal and reception, respectively. Thus, RP_SEND must know the IP address and port number of the machine where RP_RECV is located. Note that the port number is specifically for communication between RP_SEND and RP_RECV for data replication, and must be different from the one used by database server. RP_SEND uses the keyword **RP_SlAdr** in **dmconfig.ini** on the primary database to determine where to send replication data.

Syntax for **RP_SlAdr**:

```
RP_SLADR = {address[:port number]}
```

The default port number is 23001.

➲ **Example 1**

The primary database can support up to 256 slave databases. A comma or a white space can separate information for slave databases. The following uses 3 slave databases, **192.168.9.222** (with port number 5100), **Mars** (with port number 5101), and **Scorpio** (with default port number 23001):

```
RP_SLADR = 192.168.9.222:5100, Mars:5101, Scorpio
```

Since the destination to replicate data to is known, the initial time and the time interval for RP_SEND to execute database replication can be set. After deciding the schedule, RP_SEND will execute data replication periodically.

➲ **Example 2**

To set the initial time to 01/01/2000 AM 01:00 and set the time interval to one day, use the following **dmconfig.ini** settings:

```
DB_SMODE = 4                 ; Start as primary database,
                             ; and also start up RP_SEND server
RP_BTIME = 00/01/01 01:00:00 ; Initial replication time
RP_ITERV = 1-00:00:00         ; Replicate everyday
RP_SLADR = 192.168.9.222:5100, Mars:5101, Scorpio ; Replicate to 3 machines
RP_RETRY = 3                 ; Times to retry if network connection fails
RP_CLEAR = 1                 ; Clear backup journal files after sending
                    ; to the slave side
```

In the configuration file, **RP_BTime** and **RP_Iterv** are used to specify the schedule for data replication. **RP_BTime** is the starting time of sending the backup journal files to the slave side. The format of **RP_BTime** is <*year/month/day hour:minute:second*>. If it is not given, the system will set **RP_BTime** to be the starting time of primary database. The **RP_Iterv** specifies the time interval at which RP_SEND is activated. Its format is <*day-hour:minute:second*>. The default value is one day. The valid value ranges from 0 to 24855.

NOTE *These values must be set before starting the database.*

RP_ReTry sets how many times to retry if network connection fails. The value of **RP_Clear** determines whether the backup journal files should be cleared after being sent to the slave. Setting **RP_Clear** to 1 clears up the backup journal files. The default value is 0. If backup journal files are used for database replication only, clearing those files can reclaim some storage space. However, if a hardware crash occurs, the data from backup journal files is irretrievable and will not be able to be restored. In this case, a full backup against the slave database must be made to restore the primary database. Therefore, if backup journal files will be used for primary database backup as well, **RP_Clear** should be set to 0.

Setup RP_RECV and RP_APPLY Servers on a Slave

In order to start up RP_RECV and RP_APPLY servers, the start-up mode **DB_SMode** of the slave database should be set to 5.

➲ **Example 1**

A slave database can receive replication data only from one primary database; hence, it needs **RP_Primy** to specify the primary database's machine name or address.

```
RP_PRIMY = FreeBSD       ; Receive replication data from machine 'FreeBSD'
```

➲ **Example 2**

RP_RECV server uses a different port number from **DB_PtNum** to receive data from RP_SEND. For example, suppose there is a slave database located on machine NTPC, with the **RP_PtNum** keyword setting.

```
RP_PTNUM = 5100             ; Port number for RP_SEND and RP_RECV connection
DB_PTNUM = 3333             ; Port number for user database access
```

➲ **Example 3**

The primary database uses the **RP_SlAdr** keyword to retrieve the slave database's machine name or address and port number.

```
RP_SLADR = NTPC:5100       ; Slave-side machine name and address
```

In addition, the slave database has to set the **DB_BkDir** backup directory to keep backup journal files received from the primary database. After RP_APPLY has applied changes to the slave database with the backup journal files, DBMaker will automatically remove them.

SLAVE DATABASE IS READ-ONLY

Data in the slave database must be the same as the primary. It cannot accept data definitions (DDL, such as Create Table and Alter Table) and update data (such as INSERT, UPDATE, and DELETE). Thus, we can say that the slave database is read-only.

During the process of database replication, the slave database uses the backup journal files received from the primary database to restore data. The system will not lock data on the primary database during the process of restoring data. Therefore, queries to the slave database are a kind of *dirty read*. In other words, at any given point in time, the same query on both the primary and slave databases may return different values because RP_APPLY is restoring data.

START THE PRIMARY AND SLAVE DATABASES

The start mode of the primary database is different from the slave. If you want to set a database to be a primary, set the start mode to be in primary database mode. On the other hand, if you want to set a database to be a slave, set the start mode to be in slave database mode. Using the **DB_SMode** located in the **dmconfig.ini** file can specify all of the setup options. The start mode **DB_SMode** for the primary database mode is 4. The **DB_SMode** for a slave database mode is 5.

You can start the primary and slave databases separately, in no particular order.

A summary of all **dmconfig.ini** keywords, for replication, mentioned in this section follows.

➲ Example 1

In the following **dmconfig.ini** file, the primary database name is **FreeBSD**, and the slave database is **NTPC**, using the minimum settings for the primary database.

```
[MYDB] ;; Primary Config, CPU：x86, OS：FreeBSD, Name：FreeBSD
;; Database related settings
DB_DBDIR = /home/dbmaker/mydb
DB_USRBB = /home/dbmaker/mydb/MYDB.BB 2
DB_USRDB = /home/dbmaker/mydb/MYDB.DB 150
FILE1 = /home/dbmaker/mydb/data/FILE1.DB 50
FILE2 = /home/dbmaker/mydb/data/FILE2.DB 50
DB_JNFIL = JN1.JNL JN2.JNL
DB_SVADR = FreeBSD                  ; Machine name of primary database
DB_PTNUM = 3333                     ; Port number of primary database
;; journal backup server related settings
DB_BMODE = 1                        ; Database start up in BACKUP-DATA mode
DB_BKSVR = 1                        ; Start up journal backup server
DB_BKDIR = /home/dbmaker/mydb/bkdir ; Directory of backup journal files
DB_BKTIM = 00/01/01 00:00:00        ; The initial backup time
DB_BKITV = 0-12:00:00               ; Perform journal backup evey 12 hours
;; RP_SEND server related settings
DB_SMODE = 4                        ; Start as primary database
                                            ; and also start up RP_SEND server
RP_BTIME = 00/01/01 01:00:00        ; Initial replication time
RP_ITERV = 1-00:00:00               ; Replicate everyday
RP_SLADR = NTPC:5100                ; Replicate to NTPC with port no. 5100
RP_RETRY = 3                     ; retry 3 timesif network connection fails
RP_CLEAR = 1                     ; Clear backup journal files after sending
```

The following is the minimum settings for the slave database:

```
[MYDB];; Slave Config., CPU：x86, OS：MS Windows NT, Name：NTPC
;; Database related settings
DB_DBDIR = d:\dbmaker\db\mydb
DB_USRBB = d:\dbmaker\db\MYDB.BB 2
DB_USRDB = d:\dbmaker\db\MYDB.DB 150
FILE1    = d:\dbmaker\db\mydb\FILE1.DB 50
FILE2    = d:\dbmaker\db\mydb\FILE2.DB 50
DB_JNFIL = JN1.JNL JN2.JNL
DB_SVADR = NTPC
DB_PTNUM = 3333
```

```
;; RP_RECV and RP_APPLY servers related settings
DB_SMODE = 5              ; Start as primary database,
                          ; also start up RP_RECV and RP_APPLY servers
RP_PRIMY = FreeBSD        ; Receive replication data only from this machine
RP_PTNUM = 5100           ; Port number for RP_SEND and RP_RECV connection
DB_BKDIR = e:\mydb\bkdir  ; Directory of temporary backup journal files
```

EXECUTE DATABASE REPLICATION IMMEDIATELY

We mentioned that the resident server for database replication will detect whether there are changes to data and will replicate the data automatically.

➲ **Example**

Use the dmSQL tool to enter a SQL command to have DBMaker execute the database replication immediately.

```
dmSQL> Set Flush;
```

Use this command when you need to synchronize the data between a primary and slave database.

When you execute this command, the journal backup server will immediately execute the incremental backup and backup the current journal files. Then three resident servers (RP_SEND, RP_RECV, and RP_APPLY) will follow the procedure and replicate data.

VERIFY REPLICATION LOG (RP.LOG) AND ERROR LOG (ERROR.LOG)

In the process of database replication, if there are any network failures or any error messages, the system will generate a log file, ERROR.LOG, in the current database directory. Error log entries follow the following syntax:

```
yy/mm/dd hh:mm:ss Daemon name:Error number:Error message
```

➲ **Example 1**

An ERROR.LOG:

```
97/12/31 11:40:59 - RP_SEND:rc = 1503, cannot connect to server 192.72.116.130
97/12/31 11:43:36 - RP_SEND:rc = 1503, cannot connect to server 192.72.116.130
97/12/31 11:45:00 - RP_SEND:rc = 1503, cannot connect to server 192.72.116.130
```

```
97/12/31 11:50:00 - RP_SEND:rc = 1503, cannot connect to server 192.72.116.130
97/12/31 11:50:45 - RP_SEND:rc = 1503, cannot connect to server 192.72.116.130
```

Both primary and slave databases may generate the ERROR.LOG.

If the replication succeeds, the system will also generate a log file named RP.LOG. The format is:

On the primary database:

```
RP_SEND:RPID id ~ id sent at yy/mm/dd hh:mm:ss
```

On the slave database:

```
RP_RECV:RPID id ~ id sent at yy/mm/dd hh:mm:ss
RP_APPLY:RPID id ~ id applied at yy/mm/dd hh:mm:ss
```

The replication log, RP.LOG, will be on the primary and slave database servers. Every backup journal file has an ID, the RPID. In the above format, RPID indicates which journal file is being processed – RPID in the RP_SEND lines indicate which one is being sent, in the RP_RECV lines they indicate which one is being received, and in the RP_APPLY they indicate which one is being restored. Normally the RPID is the same as the incremental backup file ID.

➲ **Example 2**

RP.LOG on a primary database:

```
RP_SEND : RPID  7 ~ 10 sent to 192.72.116.130 at 97/12/16 15:36:17
RP_SEND : RPID 11 ~ 11 sent to 192.72.116.130 at 97/12/16 15:59:42
RP_SEND : RPID 12 ~ 12 sent to 192.72.116.130 at 97/12/31 11:52:28
```

➲ **Example 3**

RP.LOG on a slave database:

```
RP_RECV : RPID  7 ~ 10 received at 97/12/16 15:35:53
RP_APPLY : RPID  7 ~ 10 applied at 97/12/16 15:35:55
RP_RECV : RPID 11 ~ 11 received at 97/12/16 15:59:18
RP_APPLY : RPID 11 ~ 11 applied at 97/12/16 15:59:18
RP_RECV : RPID 12 ~ 12 received at 97/12/31 11:52:01
RP_APPLY : RPID 12 ~ 12 applied at 97/12/31 11:52:02
```

RP.LOG is used to record all actions performed by three resident servers, RP_SEND, RP_RECV, and RP_APPLY. The file is located in the database directory, **DB_DbDir**, on all replication participant databases. Database administrators should monitor these files periodically to ensure that the replication is running smoothly.

For example, if the connection to a remote database always fails, check whether the network works normally or if the remote database was engaged at the time of failure.

JServer Manager Environment Settings

To replicate a database, you need to establish the initial database environment. Use a text editor to modify the **dmconfig.ini** file directly, as shown in the previous section. In this section, DBMaker*'s* JServer Manager will be used to set up the database replication environment.

PRODUCE A FULL BACKUP

First, produce an offline full backup of the primary database, and then copy the backup to the slave database. For more information on offline full backup, refer to Section 14.5, *Offline Full Backups*.

PRIMARY DATABASE SETUP

To allow the database replication server to function, first set up the database so that incremental backups are being made. Refer to Chapter 14, *Database Recovery, Backup, and Restoration* for more information. Next, set up the environment for the primary database.

➲ **To setup for the primary database environment:**

1. Start the **JServer Manager** application on the primary database server.
2. Click on the **Setup** button in the **Start Database** window.
3. Click the Replication tab in the Start Database Advanced Settings window.
4. To enable database replication:
   - **a)** Enter slave database IP and port numbers into the **IP and Port Number of Target Database**.
   - **b)** Enter a date and time into the **Begin Time of Database Replication** time fields.
   - **c)** Enter a value into the **Times to Retry upon Failure** field

d) If desired, enable **Remove Backup Journal Files after Replication**

e) Enter the number of days, hours, minutes, and seconds between each successive database replication in the **Time Interval to Start Database Replication** time fields.

5. Click the **Save** button.

6. Click the **Cancel** button to return to the **Start Database** window.

SLAVE DATABASE SETUP

After copying a full backup of the primary database to the slave database, use JServer Manager to setup the configuration of the slave database.

1. Start the **JServer Manager** application on the slave database server.

2. Click on the **Setup** button in the **Start Database** window.

3. Click the Replication tab in the Start Database Advanced Settings window.

4. To enable database replication:

   a) Enter the IP Address of Source Database.

   b) Enter a port number for RP_RECV in the **Port Number of Receive Daemon on Target DB**.

5. Click the **Save** button.

6. Click the **Cancel** button to return to the **Start Database** window

Database Configuration File

This section is a summary of database replication related keywords used in the **dmconfig.ini** file.

PRIMARY DATABASE CONFIGURATION

Primary database keywords for data replication:

- **DB_SMode** — 1Setting **DB_SMode** to 4 means that this database will start in primary mode.

- **RP_BTime** — 1Sets the beginning time for the primary database full backup files to be sent to the slave database. The format is yr/mon/day hr:min:sec. The default value is the starting time of the primary database. For example, 97/12/31 12:00:00.

- **RP_Iterv** — 1Sets the time interval to send the backup journal files. The format is day-hr:min:sec. For instance, 1-12:00:00 means to send the backup files every one and half day. The valid range of days is 0~24855.
- **RP_ReTry** — 1Sets the number of times to retry a connection when the network fails.
- **RP_Clear** — 1Setting to clear journal backup files after sending them. The value 1 clears the files, a value of 0 does not clear them. The default value is 0. If the value is set to 1, the primary database cannot be restored if there are hardware damages unless the slave database is used to restore it.
- **RP_SlAdr** — 1This value will set the address (or machine number) and the port number of the slave database. DBMaker supports 1 to 8 slave databases for each primary database.

  Syntax for **RP_SlAdr**:

  ```
  RP_SLADR = {Address[:Port Number]}
  ```

  The primary database needs to start up the journal backup server to perform incremental backups:

- **DB_Bmode** — 1 (BACKUP-DATA) or 2 (BACKUP-DATA-AND-BLOB).
- **DB_BkSvr** — Set **DB_BkSvr** to 1 to start the journal backup server.
- **DB_BkDir** — The directory to store backup journal files.
- **DB_BkTim** — The starting time of the journal backup server. The format is <*yr/mon/day hr:min:sec*>. The default value is the starting time of the primary database.
- **DB_BkItv** — The time interval to perform incremental backups, and the format is <*day-hr:min:sec*>. For instance, 0-12:00:00 means to backup every 12 hours.

SLAVE DATABASE CONFIGURATION

The slave database keywords for data replication:

- **DB_Smode** — Setting **DB_SMode** to 5 means that this database will start in slave mode.
- **RP_Primy** — The address or machine name of the primary database.
- **RP_PtNum** — The port number used by RP_RECV. The value must be different from **DB_PtNum**, and it has to be the same as the port number specified in **RP_SlAdr** of the primary database.
- **DB_BkDir** — The directory to store the temporary backup journal files received from the primary database. The default directory is **<*database directory*>/backup**.

Database Replication Limitations

Summary of database replication limitations:

- Byte ordering must be the same on the primary and slave machines.
- Slave database is read-only.
- One primary database supports up to 256 slave databases.
- Databases for database replication cannot perform online full backups.
- Currently replication of FILE data type is not supported.
- To replicate BLOB data, i.e. columns of LONG VARCHAR or LONG VARBINARY data type, set **DB_BMode** to 2 (BACKUP-DATA-AND-BLOB) and remember to set the BACKUP BLOB ON option while creating tablespaces.

17 Performance Tuning

DBMaker is a highly tunable database system. Tuning DBMaker will increase its performance level to satisfy individual needs. This chapter presents the goals and methods used in the tuning process, and demonstrates how to diagnose a system's performance.

17.1 The Tuning Process

Before tuning DBMaker, you must define goals for improving performance. Keep in mind that some goals may conflict. You must decide which of the conflicting goals are most important.

The following outlines some of the goals for tuning DBMaker:

- Improving the performance of SQL statements.
- Improving the performance of database applications.
- Improving the performance of concurrent processing.
- Optimizing resource utilization.

After determining the goals, you are ready to begin tuning DBMaker.

Start by performing the following steps:

- Monitor database performance
- Tuning I/O.
- Tuning memory allocation.

- Tuning concurrent processing.
- Monitor database performance and compare with previous statistics

The methods used to perform tuning in each of these steps may have a negative influence on other steps. Following the order shown above can reduce this influence. After performing all of the tuning steps, monitor the performance of DBMaker to see whether the best overall performance has been achieved.

Before tuning DBMaker, make certain that the SQL statements are written efficiently and the database applications employ good design. Inefficient SQL statements or badly designed applications can have a negative influence on database performance that tuning cannot improve. To write efficient statements and applications, refer to the *SQL Command and Function Reference* and the *ODBC Programmer's Guide*.

17.2 Monitoring a Database

This section shows how to monitor information about the status of a database, including resource status, operation status, connection status, and concurrency status. This section also shows how to kill a connection.

The Monitor Tables

DBMaker stores the database status in four system catalog tables: SYSINFO, SYSUSER, SYSLOCK, and SYSWAIT.

The SYSINFO table contains database system values including total DCCA size, available DCCA size, number of maximum transactions, and the number of page buffers. It also includes statistics on system actions such as the number of active transactions, the number of started transactions, the number of lock and semaphore requests, the number of physical disk I/O, the number of journal record I/O, and more. Use this table to monitor the database system status, and use the information to tune the database.

The SYSUSER table contains connection information, including connection ID, user name, login name, login IP address, and the number of DML operations that have been executed. Use this table to monitor which users are using a database.

The SYSLOCK table contains information about locked objects, such as the ID of the locked object, lock status, lock granularity, ID of the connection locking this object, and more. Use this table to monitor which objects are being locked by which connection, and which users are locking which objects.

The SYSWAIT table contains information on the wait status of connections, including the ID of the connection that is waiting and the ID of the connection it is waiting for. Use this table to monitor the concurrency status of connections. Once a connection is waiting for those resources locked by an idle or dead connection, you can determine which connection is locking those objects from this table. Then you can kill the idle or dead connection to release the resources.

Browse these four system catalog tables in the same way ordinary tables are browsed.

➲ **Example**

SQL SELECT command used to browse the SYSUSER table:

```
dmSQL> SELECT * FROM SYSUSER;
```

Refer to Appendix B for more information on these four system catalog tables.

Killing Connections

A connection should be killed when the connection is holding resources and is idle for a long time, or when the resources are urgently required. In addition, all active connections should be killed before shutting down a database. Before killing a connection, browse the SYSUSER table to determine its connection ID.

➲ **Example 1**

To kill the connection for **Eddie**, retrieve the connection ID first:

```
dmSQL> SELECT CONNECTION_ID FROM SYSUSER WHERE USER_NAME = 'Eddie';

CONNECTION_ID
=============
352501
```

➲ **Example 2**

Then to kill the connection for Eddie use:

```
dmSQL> KILL 352501;
```

17.3 Tuning I/O

Disk I/O requires the most time in DBMaker.

To avoid disk I/O bottlenecks, perform the following:

- Determine data partitions.
- Determine journal file partitions.
- Separate journal files and data files onto different disks.
- Use raw devices.
- Pre-allocate space in an autoextend tablespace.
- Turn I/O and checkpoint daemon on.

Determining Data Partitions

You can use tablespaces to partition data instead of storing all of the data together. If tablespaces are used properly, DBMaker will have greater performance when performing space management functions or full table scans. Small tables that contain data of a similar nature can be grouped in a single tablespace, but very large tables should be placed in their own tablespace.

You can achieve speed improvement in disk I/O by using disk striping. Striping is the practice of separating consecutive disk sectors so they span several disks. This can be used to divide the data in a large table over several disks. This helps to avoid disk contention that may occur when many processes try to access the same files concurrently.

Determining Journal File Partitions

DBMaker gives the flexibility to use one or more journal files. A single journal file is easier to manage, but using multiple journal files has some advantages as well. If you run DBMaker in backup mode and use the backup server to perform incremental backups, using multiple journal files can improve the performance of incremental backups. Only full journal files will be backed up. In addition, spreading multiple journal files across different disks can increase disk I/O performance.

You may determine the size of journal files by examining the needs of transactions. However, if you run DBMaker in backup mode and perform backups according to the journal full status, the journal size will also affect the backup time interval. A larger journal file increases the interval between backups.

Separating Journal Files and Data Files

Separating the journal files and data files onto different disks will increase disk I/O performance, permitting files to be accessed concurrently to some degree. If the disks have different I/O speeds, consider which files to put on the faster disks. In general, if you run on-line transaction processing (OLTP) applications often, put the journal files on the faster disks. However, if you run applications that perform long queries, such as a decision support system, put data files into faster disks.

Using Raw Devices

If you run DBMaker on a UNIX system, construct raw device files to store DBMaker data and journal files. Since DBMaker has a good buffer mechanism, it is much faster to read/write from a raw device than a UNIX file. For more information on how to create a raw device, refer to the operating system manual or consult your system administrator. The one disadvantage of using raw devices is that DBMaker cannot extend tablespaces on them automatically; so more planning is required when using raw device files.

Pre-Allocating Autoextend Tablespaces

DBMaker supports autoextend tablespaces to simplify tablespace management. However, if you are able to estimate the required size of a tablespace, it is better to fix the size when creating the tablespace. This improves performance, as extending pages takes a lot of time. You can extend the pages of a file later by using the alter file command. Pre-allocating the size of a tablespace can also avoid disk full errors when DBMaker attempts to extend a tablespace that already occupies all available disk space.

I/O and Checkpoint Daemons

I/O Daemon

DBMaker has an I/O daemon to periodically write dirty pages from the least recently used page buffers to disk. This helps reduce the overhead incurred when swapping data pages into the page buffers, and increases performance. One configuration parameter in the **dmconfig.ini** file is used to control the I/O daemon.

DB_IOSvr—enables and disables the I/O daemon. Setting this keyword to a value of 1 enables the I/O daemon, and setting it to a value of 0 disables the I/O daemon.

➲ **Example**

A typical excerpt from the **dmconfig.ini** file:

```
[MYDB]
...
DB_IOSVR = 1
```

MYDB database has 400 (**DB_NBufs**) page buffers in DCCA. Every 10 minutes, the I/O daemon will perform the following steps:

- Scan the least recently used page buffers.
- Collect the dirty pages during scan processing.
- Write these collected dirty pages to disk.

CHECKPOINT DAEMON

DBMaker has a checkpoint daemon (based on the I/O daemon) that periodically takes a checkpoint. This helps reduce the time spent waiting for a checkpoint that occurs during a command, when a journal is full, or when starting or shutting down a database. The checkpoint daemon is actually a sub-function of the I/O daemon, which can perform I/O alone, checkpoints alone, or both together. There is one keyword for use in the **dmconfig.ini** file, which is used to control the checkpoint daemon.

DB_ChTim—specifies the first time a checkpoint daemon should run. The format for this keyword is yyyy/mm/dd hh:mm:ss.

To turn on the checkpoint daemon, turn on the I/O daemon using the **DB_IOSrv** keyword. If the I/O daemon is activated without setting **DB_ChTim**, it will automatically take a checkpoint every hour by default after the database starts successfully.

➲ **Example 2**

To start checkpoint daemon and stop the I/O daemon in the **dmconfig.ini**:

```
[MYDB]
…
DB_IOSVR = 1                    ; may enable I/O or checkpoint daemon
DB_CHTIM = 2000/1/1 00:00:00    ; the first time the daemon should run
```

In fact, the I/O and checkpoint daemon will expend some I/O resources. After starting the database server, any error messages generated by the I/O and checkpoint daemon are written to the file **ERROR.LOG**.

17.4 Tuning Memory Allocation

DBMaker stores information temporarily in memory buffers and permanently on disk. Since it takes much less time to retrieve data from memory than disk, performance will increase if data can be obtained from the memory buffers. The size of each of DBMaker’s memory structures will affect the performance of a database. However, performance becomes an issue only if there is not enough memory.

This section focuses on tuning the memory usage for a database. It includes information on how to calculate the required DCCA size, and how to monitor and allocate enough memory for the page buffers, journal buffers and system control area.

➲ **To achieve the best performance, follow the steps in the order shown:**

1. Tune the operating system.
2. Tune the DCCA memory size.
3. Tune the page buffers.
4. Tune the journal buffers.
5. Tune the SCA.

DBMaker's memory requirement varies according to the applications in use; tune memory allocation after tuning application programs and SQL statements.

Tuning an Operating System

The operating system should be tuned to reduce memory swapping and ensure that the system runs smooth and efficiently.

Memory swapping between physical memory and the virtual memory file on disk takes a significant amount of time. It is important to have enough physical memory for running processes. Measure the status of an operating system with the operating system utilities. An extremely high page-swapping rate indicates that the amount of physical memory in a system is not large enough. If this is the case, remove any unnecessary processes or add more physical memory to the system.

Tuning DCCA Memory

The *Database Communication and Control Area (DCCA)* is a group of shared memory allocated by DBMaker servers. Every time DBMaker is started, it allocates and initializes the DCCA.

The UNIX client/server model of DBMaker allocates the DCCA from the UNIX shared memory pool. Ensure that the size of the DCCA is not larger than the maximum-shared memory size permitted by the operating system. If the requested size for the DCCA is larger than the operating system limit, refer to the operating system administration manual for information on how to increase the maximum size of shared memory.

CONFIGURING THE DCCA

The DCCA contains the process communication control blocks, concurrency control blocks, and the cache buffers for data pages, journal blocks, and catalogs. DBMaker maintains the concurrency control blocks and communication status of each DBMaker process in the DCCA. Each DBMaker process accesses the same disk data through the cache buffers in the DCCA.

Setting the appropriate parameters in **dmconfig.ini** before starting the database configures the size of each of the DCCA components.

➲ **Example 1**

A sample configuration for the DCCA in the **dmconfig.ini** file:

```
DB_NBUFS = 200
DB_NJNLB = 50
DB_SCASZ = 50
```

DB_NBufs specifies the number of data page buffers (4096 bytes per buffer), **DB_NJnlB** specifies the number of journal block buffers (4096 bytes per buffer), and **DB_ScaSz** specifies the size of the SCA in pages (4096 bytes per page). DBMaker reads these DCCA parameters only when starting a database. To adjust the parameters, terminate the database, modify the values in the **dmconfig.ini** file, and restart the database. For more information on setting these parameters, refer to Appendix A.

The total memory allocation for the DCCA is the sum of the size of **DB_NBufs**, **DB_NJnlB** and **DB_ScaSz**.

➲ **Example 2**

To calculate the total size of the DCCA:

```
size of DCCA  = (200 + 50 + 50) * 4 KB
                           = 1200 KB
```

ALLOCATING SUFFICIENT DCCA PHYSICAL MEMORY

The DCCA is the resource most frequently accessed by DBMaker processes. It is important to ensure there is enough physical memory to prevent the operating system from swapping the DCCA to disk too often or it will seriously degrade the performance of a database. The page-swapping rate can be measured by using operating system utilities.

➲ **Example**

To determine the size of memory allocated for the DCCA from the system table SYSINFO:

```
dmSQL> select INFO, VALUE from SYSINFO where INFO = DCCA_SIZE
                                          or INFO = FREE_DCCA_SIZE;

            INFO                          VALUE
============================= ===============================
DCCA_SIZE                     1228800
FREE_DCCA_SIZE                 189024
```

DCCA_SIZE—the memory size, in bytes, of the DCCA.

FREE_DCCA_SIZE—the size, in bytes, of free memory remaining in the DCCA.

The free memory in the DCCA is reserved for use by dynamic control blocks, such as lock control blocks.

Usually a larger number of buffers is better for system performance. However, if the DCCA is too large to fit in physical memory, the system performance will degrade. Therefore, it is important to allocate enough memory for the DCCA but still fit the DCCA in physical memory.

Tuning Page Buffer Cache

DBMaker uses the shared memory pool for the data page buffer cache. The buffer cache allows DBMaker to speed up data access and concurrency control. DBMaker automatically configures the number of page buffers by default. Setting the **dmconfig.ini** keyword **DB_Nbufs** to zero allows DBMaker to automatically set the number of page buffers. DBMaker can dynamically adjust the number of page buffers on systems that allow DBMaker to detect physical memory usage. The number will be no less than 500 pages on Windows 95/98, or no less than 2000 pages for Windows NT/2000/ or Unix. If DBMaker cannot detect the system's physical memory usage, it will allocate the minimum amount.

Adjusting the size of the page buffers will have the greatest effect on performance. The next sections show how to monitor the buffer cache performance and calculate the buffer hit ratios.

➲ **To improve buffer cache performance:**

1. Update statistics on schema objects.
2. Set NOCACHE on large tables.
3. Reorganize data in poorly clustered indexes.
4. Enlarge cache buffers.
5. Reduce the effect of checkpoints.

MONITORING PAGE BUFFER CACHE PERFORMANCE

DBMaker places buffer cache access statistics in the SYSINFO system table.

➲ **Example**

To get buffer cache values use the following SQL statements:

```
dmSQL> select INFO, VALUE from SYSINFO where INFO = 'NUM_PAGE_BUF';

              INFO                               VALUE
================================ ================================
NUM_PAGE_BUF                     4000

1 rows selected

dmSQL> select INFO, VALUE from SYSINFO where INFO = 'NUM_PHYSICAL_READ'
    2>                                       or INFO = 'NUM_LOGICAL_READ'
    3>                                       or INFO = 'NUM_PHYSICAL_WRITE'
    4>                                       or INFO = 'NUM_LOGICAL_WRITE';

              INFO                               VALUE
================================ ================================
NUM_PHYSICAL_READ                64
NUM_PHYSICAL_WRITE               1
NUM_LOGICAL_READ                 509
NUM_LOGICAL_WRITE                0
```

```
4 rows selected
```

NUM_PAGE_BUF—number of pages used for data buffer cache.

NUM_PHYSICAL_READ—number of pages read from disk.

NUM_LOGICAL_READ—number of pages read from the buffer cache.

NUM_PHYSICAL_WRITE—number of pages written to disk.

NUM_LOGICAL_WRITE—number of pages written to the buffer cache.

Calculate the page buffer read/write hit ratio with the following formulas:

$$\text{read hit ratio} = 1 - \left(\frac{\text{NUM\_PHYSICAL\_READ}}{\text{NUM\_LOGICAL\_READ}}\right)$$

$$\text{write hit ratio} = 1 - \left(\frac{\text{NUM\_PHYSICAL\_WRITE}}{\text{NUM\_LOGICAL\_WRITE}}\right)$$

Using the example above, can calculate the read/write hit ratio:

$$\begin{aligned}\text{read hit ratio} &= 1 - \left(\frac{13207}{331595}\right)\\ &= 0.960\\ &= 96.0\%\end{aligned}$$

$$\begin{aligned}\text{write hit ratio} &= 1 - \left(\frac{7361}{127423}\right)\\ &= 0.942\\ &= 94.2\%\end{aligned}$$

Based on the read/write hit ratio, determine how to improve the buffer cache performance. If the hit ratio is too low, tune DBMaker with the methods described in the following subsections.

If the hit ratio is always high, for example higher than 99%, the cache is probably large enough to hold all of the most frequently used pages. In this case, try to reduce the cache size to reserve memory for applications. To ensure good performance, monitor the cache performance before and after making the modifications.

Statistics Values are Outdated

If the read/write hit ratio is too low, it may be that the statistics values of schema objects (tables, indexes, columns) are out of date. The wrong statistics may cause the DBMaker optimizer to use an inefficient plan for SQL statement. If users have inserted large amounts of data into the database after the last time the statistics values were updated, update the values again.

➲ **Example 1**

To update the statistics values for all schema objects:

```
dmSQL> update statistics;
```

If a database is extremely large, it will take a lot of time to update statistical values for all of the schema objects. An alternative method is to update statistics on specific schema objects that have been modified since the last update, and set the sampling rate.

➲ **Example 2**

To update specific schema objects:

```
dmSQL> update statistics tabel1, table2, user1.table3 sample = 30;
```

After successfully updating the statistical values for schema objects, monitor the performance of the page buffer cache with the method specified in *Monitoring Page Buffer Cache Performance*.

Swap Out Cache

DBMaker determines which page buffers to swap with the *Least Recently Used (LRU)* rule. This keeps the most frequently accessed pages in the page buffers and swaps pages that are used less frequently. However, if a large table is browsed all page buffers may be swapped out just to perform one table scan.

For example, in a database with 200 page buffers, if a table with 250 pages is browsed, DBMaker might read all 250 pages into the page buffers and discard the 200 most frequently used pages. In the worst case, DBMaker must read 200 pages from disk when accessing other data after a full table scan. However, if the table cache mode is set to NOCACHE, DBMaker will place the retrieved pages at the end of the LRU chain when a full table scan is performed. Therefore, 199 of the 200 most frequently used pages are still kept in the buffer cache.

Normally the tables with page numbers that exceed the page buffers should be set to NOCACHE. Tables that are not used frequently or with page numbers close to the number of page buffers should also be set to NOCACHE.

➲ **Example 1**

To determine the number of pages and cache mode for a table:

```
dmSQL> select TABLE_OWNER, TABLE_NAME, NUM_PAGE, CACHEMODE from SYSTEM.SYSTABLE
where TABLE_OWNER != 'SYSTEM';

TABLE_OWNER     TABLE_NAME      NUM_PAGE   CACHEMODE
=========== ================== =========== =========
       BOSS             salary           5 T
        MIS              asset          45 T
        MIS         department           3 T
        MIS           employee          29 T
        MIS           worktime         450 T
      TRADE           customer         350 T
      TRADE          inventory         167 T
      TRADE              order         112 T
      TRADE        transaction        1345 F

9 rows selected
```

NUM_PAGE—the number of pages in a table.

CACHEMODE—cache mode of full table scan, 'T' means table scan is cacheable, and 'F' means table scan is non-cacheable.

In the above sample, the table **TRADE.transaction** is already set to NOCACHE. The other tables still are cacheable. If there are 200 page buffers, the **MIS.worktime** and **TRADE.customer** tables should be set to NOCACHE, and the **TRADE.order** and **TRADE.inventory** tables should be set to NOCACHE if they are rarely used.

➲ **Example 2**

To set the cache mode for a table to NOCACHE:

```
dmSQL> alter table MIS.worktime set nocache on;
```

If there are no valid indexes for a table, or the predicate in a query references non-indexed columns, DBMaker may also perform a full table scan. To prevent this, try to write SQL statements as efficiently as possible, and make use of indexed columns when possible.

Poor Clustering of Records

When fetching many records that must be ordered by an index key, or when the predicate references an indexed column, index clustering becomes an important factor that affects the buffer cache performance.

➲ **Example 1:**

To select all columns from the customer table and sort it using the **custid** primary key:

```
dmSQL> SELECT * FROM customer ORDER BY custid;
```

Suppose there are 3500 records in the **customer** table distributed over 350 pages, and there are 200 page buffers in the system. If the records are clustered using **custid** and the clustering is very good (arranged sequentially on all pages), DBMaker only needs to read 350 pages from disk. On the other hand, if the clustering is poor (no sequential records on the same page), DBMaker may have to read 3500 pages from disk in the worst case (every record needs a disk read)! To determine the state of an index cluster, update statistics on the table first.

➲ **Example 2**

To build an index called **custid_index** on the **custid** column for the **customer** table:

```
dmSQL> select CLSTR_COUNT from SYSTEM.SYSINDEX
                          where TABLE_OWNER = 'TRADE'
                          and TABLE_NAME = 'customer'
                          and INDEX_NAME = 'custid_index';
```

Result:

```
CLSTR_COUNT
===========
```

```
      385

1 rows selected
```

CLSTR_COUNT—cluster count, the number of data pages that will be fetched by a fully indexed scan with few buffers. DBMaker performs 385 page reads from disk at most, when the full customer table is scanned and results ordered by the **custid** column.

➲ **Example 3**

To fetch the number of pages and rows:

```
dmSQL> select NUM_PAGE,NUM_ROW from SYSTEM.SYSTABLE
                               where TABLE_OWNER = 'TRADE'
                               and TABLE_NAME = 'customer';
```

Result:

```
NUM_PAGE    NUM_ROW
=========== ===========
   350        4375

1 rows selected
```

NUM_PAGE—the number of pages allocated by a table.

NUM_ROW—the number of records in a table.

With CLSTR_COUNT, NUM_PAGE and NUM_ROW, estimate the clustering factor with the following formula:

$$\text{clustering factor} = \frac{(\text{CLSTR\_COUNT - NUM\_PAGE})}{\text{NUM\_ROW}}$$

In the above example, the clustering factor is 1.7%.

$$\begin{aligned}\text{clustering factor} &= \frac{(385 - 350)}{9375} \\ &= 0.0017 \\ &= 1.7\%\end{aligned}$$

The clustering factor will be between 0 and 100%. In cases where CLSTR_COUNT is only a little less than NUM_PAGE, it can be treated as zero. If the clustering factor is zero, it means that the data is fully clustered for the index. If the clustering factor is too high, for example larger than 20% (what determines a high rate depends on the table size, average record size, etc.), the index has poor clustering. When DBMaker finds an index that has poor clustering, the DBMaker optimizer may use a full table scan when an SQL statement executed, even if an index scan seems more appropriate.

➲ **To improve poor clustering for a frequently used index:**

1. Unload all data from the table (ordered by the index).
2. Rearrange the unloaded data by order.
3. Drop indexes on the table.
4. Delete all data in the table.
5. Reload the data into the table.
6. Recreate indexes on the table.

After data reloading, the index should be fully clustered. Note however, a table can only be clustered with one index. If one table has many indexes, maintain clustering on the most important index. Usually, the most important index is the primary key. Since unloading and reloading data takes a great deal of time and storage, tune index clustering only on the tables that are very large and frequently browsed.

Low Data Page Buffers

If allocated data page buffers are not enough for database access, add page buffers to the DCCA.

➲ **To modify the number of page buffers:**

1. Terminate the database server.
2. Reset **DB_NBufs** in **dmconfig.ini** to a larger value.
3. Restart the database.

After successfully enlarging the data buffers, run the database for a period and then monitor the buffer cache performance again. If the buffer hit ratio has gone up, adding buffer pages has resulted in a performance improvement. If not, add more pages to the buffer cache or check for other reasons why system performance may be reduced.

CHECKPOINTS OCCURRING TOO OFTEN

If the write hit ratio is much lower than the read hit ratio, the checkpoints may be processed too often.

When a checkpoint is processed, DBMaker will write all dirty page buffers to disk. Since checkpoints require a lot of CPU time, specify the checkpoint daemon to perform a checkpoint on a regular schedule. Another advantage of performing checkpoints periodically is to reduce the recovery time taken by DBMaker to start a database after a system crash.

Except when the checkpoint daemon makes checkpoints regularly, DBMaker will perform checkpoints automatically when running out of free journal space in NON-BACKUP mode or when an incremental backup is performed in BACKUP mode. To increase the time interval between these kinds of checkpoints, enlarge the journal size.

➲ **Example 1**

To determine how many checkpoints have been processed:

```
dmSQL> select INFO, VALUE from SYSINFO where INFO = 'NUM_CHECKPOINT';

              INFO                                VALUE
================================ ================================
NUM_CHECKPOINT                   26

1 rows selected
```

RE-MONITOR THE BUFFER CACHE PERFORMANCE

After tuning a system with the above methods, monitor the cache buffer performance.

➲ **To monitor cache buffer performance:**

1. Run the database for a period to ensure the information in the database is in a stable state.

2. Reset the statistics values in the SYSINFO system table with the following:

```
dmSQL> set SYSINFO clear;
```

3. Run the database for a period time.
4. Get the read/write counter from the SYSINFO table and check the hit ratio.

Tuning Journal Buffers

The journal buffers store the most recently used journal blocks. With enough journal buffers, the time required to write journal blocks to disk when updating data and reading journal blocks from disk when rolling back transactions is reduced.

If you seldom run a long transaction that modifies (inserts, deletes, updates) many records, you may skip this section. Otherwise, should determine whether there are sufficient journal buffers for the system. The optimum number of journal buffers is the sum of journal blocks needed by the longest running transactions at the same time.

➲ **To estimate the number of journal buffers, perform the following:**

1. Make sure there is only one active user in the database.
2. Clear the counters in the SYSINFO table with the following command:

```
dmSQL> set SYSINFO clear;
```

3. Run the transaction that will update the most records.
4. Run the following SQL statement to determine the number of used journal blocks:

```
dmSQL> select INFO, VALUE from SYSINFO where INFO =
'NUM_JNL_BLK_WRITE';


              INFO                                VALUE

================================ ================================

NUM_JNL_BLK_WRITE                626
```

```
1 rows selected
```

NOTE NUM_JNL_BLK_WRITE—*the blocks used in this transaction. The journal block size used in this example is 512 bytes. In the above example, you need approximately 41 journal buffer pages (1 page=4KB).*

Another measurement that can be used to determine the journal buffer utilization is the journal buffer flush rate. The journal buffer flush rate is the percentage of journal buffers flushed to disk when DBMaker writes to the journal. If the journal buffer flush rate is too high (for example, more than 50%), increase the number of journal buffers.

➲ **Example 1**

To calculate the journal buffer flush rate:

```
dmSQL> select INFO, VALUE from SYSINFO where INFO = 'NUM_JNL_BLK_WRITE'
    2>                                        or INFO = 'NUM_JNL_FRC_WRITE';

              INFO                                VALUE
================================ ================================
NUM_JNL_BLK_WRITE                41438
NUM_JNL_FRC_WRITE                159

2 rows selected
```

NUM_JNL_BLK_WRITE—number of journal blocks written to the buffer.

NUM_JNL_FRC_WRITE—number of times a forced write of the journal buffers to disk occurred.

Suppose **DB_NJnlB** is set to 50 pages (i.e. there are 400 journal buffers). In the example below, the journal flush rate (0.65) is a little too high. Add journal buffers to improve the journal buffer performance.

$$\begin{aligned}\text{journal flush rate} &= \frac{(\text{NUM\_JNL\_BLK\_WRITE}/\text{NUM\_JNL\_FRC\_WRITE})}{(\text{DB\_NJNLB}\times 8)} \\ &= \frac{(41438/159)}{(50\times 8)} \\ &= 0.65\end{aligned}$$

Tuning the System Control Area (SCA)

Cache buffers and some control blocks, such as session and transaction information, have a fixed size, and are pre-allocated from the DCCA when a database is started. However, some concurrency control blocks are allocated dynamically from the DCCA while the database is running, their size is specified by **DB_ScaSz**.

If a database application gets the error message "*database request shared memory exceeds database startup setting*", it means that DBMaker cannot dynamically allocate memory from the SCA area. Usually, this error is due to a long transaction using too many locks. If this situation happens often, solve it with the methods illustrated below.

AVOID LONG TRANSACTIONS

A long transaction will occupy many lock control blocks and journal blocks. If there is a long transaction in progress when the above error occurs, analyze whether the transaction can be divided into multiple small transactions.

AVOID EXCESSIVE LOCKS ON LARGE TABLES

Selecting many records from a large table using an index scan requires many lock resources. To decrease the amount of lock resources used by the transaction, escalate the lock mode before performing the table scan.

For example, if the table's default lock mode is row, escalate the default lock mode to page or table. Although this will reduce the resources used, it will also sacrifice concurrency to some degree.

INCREASE THE SCA SIZE

If both of the above conditions have not occurred, increase the size of the SCA. Reset the value of **DB_ScaSz** in **dmconfig.ini** to a larger value and then restart the database.

Tuning the Catalog Cache

DBMaker stores the catalog cache in the SCA. If schema objects are seldom modified, turn on the data dictionary turbo mode by setting **DB_Turbo**=1 in the **dmconfig.ini** file. When turbo mode is on, DBMaker will extend the lifetime of the catalog cache. This can improve the performance of on-line transaction processing (OLTP) programs.

17.5 Tuning Concurrent Processing

Resource contention occurs in a multi-user database system when more than one process tries to access the same database resources simultaneously. This can also lead to a situation known as a deadlock, which occurs when two or more processes wait for each other. Resource contention causes processes to wait for access to a database resource, reducing system performance.

DBMaker provides the following methods to detect and reduce resource contention:

- Reducing lock contention.
- Limiting the number of processes.

Reducing Lock Contention

When accessing data in a database, DBMaker processes will lock the target objects (records, pages, tables) automatically. When two processes want to lock the same object, one must wait. If more than two processes wait for the other processes to release the lock, a deadlock occurs. When a deadlock occurs, DBMaker will sacrifice the last transaction that helped cause the deadlock by rolling it back. Deadlock reduces system performance. Monitor lock statistics to avoid a deadlock in DBMaker.

➲ **Example 1**

To view deadlock statistics:

```
dmSQL> select INFO, VALUE from SYSINFO where INFO = 'NUM_LOCK_REQUEST'
    2>                                       or INFO = 'NUM_DEADLOCK'
    3>                                       or INFO = 'NUM_STARTED_TRANX';
```

```
           INFO                              VALUE
================================ ================================
NUM_STARTED_TRANX                33
NUM_LOCK_REQUEST                 173
NUM_DEADLOCK                     0

3 rows selected
```

NUM_LOCK_REQUEST—the number of times a lock was requested.

NUM_DEADLOCK—the number of times deadlock occurred.

NUM_STARTED_TRANX—the number of transactions that have been issued.

In the above example, on average one transaction is in deadlock per 51 (9287/181) transactions and one transaction requests approximately 83 (772967/9287) locks.

If the deadlock frequency is high, examine the schema design, SQL statements, and applications. Setting the table default lock mode lower, such as ROW lock, could reduce the lock contention, but it will require more lock resources.

Another method is to use the browse mode to read a table on a long query if the data does not need to remain consistent after the point in time that it was retrieved. This is useful when viewing the data or performing calculations using the data while not performing any updates. It provides a snapshot of the requested data at a particular point in time, but with the benefit of increased concurrency and fewer lock resources consumed, because locks are freed as soon as the data is read.

Limiting the Number of Processes

DBMaker allows up to 1200 simultaneous session connections to a server. If server resources, (such as memory, CPU power) are not sufficient, limit the maximum number of connections to avoid resource contention. The configuration parameter **DB_MaxCo** affects the maximum number of connections in the database.

When a database is initially created, the journal file is formatted for a specific number of connections. The journal file needs to be able to preserve a transaction information array for each connection. The number of connections available according to the journal file is also known as the *hard connection number*. This value is determined by the value of **DB_MaxCo** when the database is created. The *hard connection number* has a minimum value of 240, a maximum value of 1200, and must be a multiple of 40. If **DB_MaxCo** is set to a value that is not a multiple of 40, then the *hard connection number* is rounded up to a value that is a multiple of 40. The hard connection number is a limitation of the journal file, therefore, to change it the database must be started in new journal mode.

The hard connection number does not directly affect the size of the DCCA. This is determined by a value known as the *soft connection number*. The soft connection number is exactly the value of **DB_MaxCo**. The soft connection number determines the number of connections that the DCCA will support, and consequently the memory usage of the DCCA. The soft connection may be any value less than or equal to the hard connection value. To change the soft connection number, restart the database normally after changing **DB_MaxCo**.

➲ **Example 1**

In the following configuration file, the hard connection number for **DB1** is 240. For database **DB2**, it is 1120.

```
[DB1]
DB_MaxCo = 50     ;; the hard connection number = 240
                  ;; the soft connection number = 50
[DB2]
DB_MaxCo = 1100   ;; the hard connection number is 1120
                  ;; the soft connection number = 1100
```

➲ **Example 2**

After starting the database successfully, the new hard connection number for **DB1** becomes 280.

```
[DB1]
DB_SMode = 2             ;; startup with new journal file
DB_MaxCo = 280           ;; the new hard connection number = 280
```

Example 3

Assuming **DB2** has already been created as in example 1, the following entry in the **dmconfig.ini** file will result in a hard connection number of 1120 and a soft connection number of 20.

```
[DB2]
DB_SMode = 1                ;; normal start
DB_MaxCo = 20               ;; the new soft connection number = 20
```

18 Query Optimization

This chapter is an introduction to the query optimizer for DBMaker. The query optimizer makes a query on SQL commands much faster and efficient by choosing the best internal execution method.

The contents in this chapter cover the following topics:

- What is query optimization and why do we need it? When you understand the goal of query optimization, you will find the role it plays in a SQL query.
- What is the Query Execution Plan (QEP) and how do you read a QEP? When you know the QEP, you will learn how DBMaker executes a SQL query command.
- How does the query optimizer operate? When you understand the way the query optimizer searches for a QEP, you can help it to find a more efficient QEP by rewriting an equivalent SQL query.
- What is the cost function? When you know how much time it takes for an operation in QEP, you will learn how the query optimizer chooses a proper operation. Use some commands provided by DBMaker to help the query optimizer find a better operation.
- What are the statistics values and where have these values been used? When you understand the usage of statistics values in query optimization, you will see the reason why the query optimizer chooses a particular execution plan.

- How can the execution speed of a query be accelerated? When you know how to write an efficient query, you can enhance the execution efficiency by rewriting the query command.

18.1 What is Query Optimization?

Data Manipulation Language (DML) commands, such as SELECT, INSERT, DELETE, and UPDATE, are a very important stage in query optimization. DBMaker may have several execution methods for one SQL query. The goal of query optimization is to find the most efficient execution plan. The main job of the query optimizer is to decide each operation, and the order in which it operates.

➲ **To find the most efficient operation:**

1. *Read the data from a table* -- can be read with a sequential or index scan.
2. *Join tables* -- tables can be joined with a nested loop or a sort merge join.
3. *Sort* -- when is a sort needed, before an operation or after it - or can sorting be avoided with an alternative solution?

The query optimizer must estimate the number of rows affected by outer joins in order to optimize the sequence of the tables to be joined. Some database users get familiar enough with the data characteristics that they become more efficient than the query optimizer at executing a query.

Query optimizer for DBMaker will estimate all possible execution plans for each plan, compute the number of rows, how many disk page I/O required, and CPU time it takes for a single table. From the above factors mentioned, find a plan with the lowest cost.

When DBMaker seeks for a query execution plan, it will consider some major operations:

- ***Table scan*** -- or called sequential scan, which means receiving each row from data pages from a database in sequential order.
- ***Index scans*** -- the order to retrieve data is referenced by the address of the data page that is pointed to by the index page.

- ***Nested join*** -- compare two or more tables, row by row, to achieve the goal of merging.
- ***Merge join*** -- sort two tables and compare row by row to achieve merging.
- ***Sort*** -- executes sort.
- ***Temporary table*** -- in the process of query execution, establish a temporary table.

➲ **Example 1**

Using the ORDER BY clause to sort:

```
dmSQL> SELECT * FROM t1, t2 WHERE t2.c2=3 AND t1.c1=t2.c1 ORDER BY
t1.c2;
```

➲ **Example 2**

Query Execution Plan 1:

```
sort t1.c2
     merge join t1.c1=t2.c1
        index scan t1 on idx1(c1)
        sort t2.c1
           table scan t2, filter t2.c2=3
```

➲ **Example 3**

Query Execution Plan 2:

```
nested join
     index scan t1 on idx2(c2)
     table scan t2, filter t2.c2=3 and t1.c1=t2.c1
```

18.2 How Does the Optimizer Operate?

When the optimizer for DBMaker performs the optimization process for a query, it will use the following rules:

- Analyze the query, then cut the WHERE predicate into several factors.

- Search all possible execution sequences and the join sequence.
- Decide whether to use a nested or sort merge join.
- Decide whether to use a table or index scan.
- Decide the sort order.

Input of Optimizer

The precision of estimation is the critical factor deciding whether the optimizer will be successful or not. However, there is only finite information for the optimizer to estimate the time required for an operation, only a small part compared with the actual execution time. All the information needed by the optimizer comes from the system catalog tables. To make sure that this information is useful and not out of date, use the UPDATE STATISTICS command. Refer to section 18.4, *Statistics* for more information.

The data listed in system catalog tables includes:

- Number of rows in a table.
- Number of data pages used by a table.
- Average bytes for a row in a table.
- Average bytes that a column uses.
- The distinct value of each index column.
- The second maximum and minimum value for each column; the reason that we do not choose the maximum and minimum value is to avoid special large and small values from affecting precision.
- Number of index scan pages occupied by the B-tree index.
- Number of levels (height) in the B-tree index.
- Number of leaf pages in the B-tree index.
- Cluster count for the B-tree index.

The premise for the optimizer to use this information is that that the data values are distributed uniformly. If the distribution of data is askew, not uniform, the optimizer will choose a poor plan.

Factors

The first job of the optimizer is to examine all expressions in the WHERE predicate. We can decompose these expressions into several smaller independent expressions called factors.

➲ **Example 1**

The optimizer will decompose the WHERE predicate into two factors **t1.c1=t2.c1** and **t1.c2=3**:

```
dmSQL> select * from t1, t2 where t1.c1=t2.c1 and t1.c2=3;
```

➲ **Example 2**

Using the WHERE predicate for one factor **t1.c1=t2.c1** or **t1.c2=3**:

```
dmSQL> select * from t1, t2 where t1.c1=t2.c1 or t1.c1=3;
```

➲ **Example 3**

Using the WHERE predicate for two factors **t1.c1=t2.c1**, and **t1.c2=3** or **t2.c2=5**:

```
dmSQL> select * from t1, t2 where t1.c1=t2.c1 and (t1.c2=3 or t2.c2=5);
```

➲ **Example 4**

Using the WHERE predicate for one factor **t1.c1=t2.c1** and **t1.c2=3** or **t2.c2=5**:

```
dmSQL> select * from t1, t2 where t1.c1=t2.c1 and t1.c2=3 or t2.c2=5;
```

From the above example, we find that when the expression contains the binary operation "**and**" that it can be divided into different factors. However, when it contains the binary operation "**or**", the decomposition is not allowed.

To find the factors, the optimizer needs to estimate the selectivity of each factor. The selectivity is the ratio of data filtered by each factor, its value is between 0 and 1. Table1 contains 100 rows.

➲ **Example 5**

Using 5 rows for query on table **t1**, **t1.c1=3** is 5/100, that is 0.05

```
dmSQL> select * from t1 where t1.c1=3;
```

If there is more than one factor in an expression, then the selectivity of this expression is the product of these factors because they are independent of each other.

Join Sequence

The join sequence is the access order of the original table to be merged. Different join sequences will produce different execution sequences and different execution times. However, no matter how we execute, we will always retrieve the correct result.

➲ **Example 1**

```
dmSQL> select * from t1, t2 where t1.c1=t2.c1;
```

Query Execution Plan 1:

```
nested join
     table scan t1
     table scan t2, filter t1.c1=t2.c1
```

Query Execution Plan 2:

```
nested join
     table scan t2
     index scan t1 on t1(c1), filter t1.c1=t2.c1
```

➲ **Example 2**

```
dmSQL> select * from t1, t2, t3 where t1.c1=t2.c1 and t2.c1=t3.c1;
```

Query Execution result, there will be 3! (=6) join sequences:

```
(t1, t2), t3
(t1, t3), t2
(t2, t1), t3
(t2, t3), t1
(t3, t1), t2
(t3, t2), t1
```

DBMaker will search all these join sequences, then compute the cost, and choose the best one.

Nested Join and Merge Join

Nested and merge joins are the two join methods supported by DBMaker.

- Nested join uses nested loop over two layers to accomplish the join purpose. The analysis algorithm for its time complexity is "n^2".
- Merge join will sort two tables in advance, and then merge these two tables with the sorted order row by row. The time complexity for sort is "n x log(n)". The data that has already been sorted has the time complexity to perform a join of "n". Sort merge join can only be used for equal joins.

From the view of the time complexity, merge join is better than a nested join. However, there are still exceptions, for example, the difference in the number of rows in two tables is great. Regardless, the optimizer will decide the best way to perform a join with cost functions and statistical values.

Table Scan and Index scan

Table scans acquire rows from a table sequentially, row by row. For instance, if a user wants to find all rows in a table that match the condition **age > 50**, then it will receive each row from each data page, then compare each row with the matched condition to retrieve the desired one.

Another scan type is called an index scan, which builds the index on columns in a table, then finds all required data by referencing the index. The index method used by DBMaker is a B-tree, and the precondition to use an index scan is to build an index on the column for use with the predicate.

Sort

Another important operation of the query optimizer is to determine how to sort, before or after a join or to try avoiding a sort.

➲ **Example:**

Creating a sort:

```
dmSQL> select * from t1, t2 where t1.c1=t2.c1 order by t1.c2;
```

Query Execution Plan 1, the optimizer performs sorting after merging:

```
sort t1.c2
      merge join t1.c1=t2.c1
           index scan t1 on idx1(c1)
           sort t2.c1
               table scan t2
```

Query Execution Plan 2, the optimizer performs sorting before merging

```
nested join
     index scan t1 on idx2(c2)
     table scan t2, filter t1.c1=t2.c1
```

18.3 Time Cost of a Query

Time to read data from disk and time to compare the column values are two major parts in performing a query.

CPU Cost

The database server must process data in memory. It has to read a row into memory, and then use a filter expression for testing. It needs to load data from two tables into memory first and then test their join condition. In addition, the database server has to collect data for the selected columns from each row. More time is required for sorting when wild cards such as "**like**" and "**match**" are used in SQL statements.

I/O Cost

It takes much more time to read a row from the disk than to check a row in memory, so one of the main purposes of the optimizer is to reduce disk I/O.

The basic unit of disk storage that a database server processes is called a page. A page is composed of clustered blocks, and the size of a page is related to the database server. The size is 4096 bytes for DBMaker. The row capacity of a page is related to the size of a row. There are usually 10 to 100 rows in a data page. The entity of an index page contains a key value and a four-byte pointer. There are usually 100 to 1000 entries in an index page.

The database server needs memory to store copies of disk page reads from the disk for processing. Due to memory limitations, some pages might be reread. The memory is called a page buffer. If the page needed happens to be in the page buffer, the server will not read the row from the disk anymore, and it will increase performance. The database server and the operating system decide the size of a disk page and the number of page buffers. The real cost to read a page is a variable and hard to estimate.

Buffers are combinations of the following factors:

- **Buffers** – It is possible for the target page to be in the page buffer. The access cost can be almost omitted in this condition.
- **Contention** – If there is more than one application program attempting to use hardware devices, such as a disk, requests from the database server will be delayed.
- **Seek time** – This is the most time-consuming operation for a disk. It means the elapsed time to move the read/write head to the location of the desired data. It is affected by the speed of the disk and the initial position of the disk read/write head. The variation also depends largely on seek time.
- **Latency time** – Also known as the rotation delay time, is related to the speed of the disk and location of the read/write head.

Cost of Table Scan

The time spent to scan all data in a table. Whether there are predicates in the query or not, it needs to compare all data contained in the pages. The cost of a table scan equals to the number of data pages.

Cost of Index Scan

Index scans read data through B-tree index pages. There are two kinds of index scans: one will read data pages referenced by the B-tree, and the other will read data directly from the index leaf; known as a leaf scan.

➲ **Example 1**

Using table **t1** with columns **c1** and **c2**, and scanning an index built on **c1**:

```
dmSQL> select * from t1 where c1 > 0;
```

Alternatively use:

```
dmSQL> select c1 from t1 where c1 > 0;
```

We can use a leaf scan and there will be no need to read data from the data pages because the leaf pages contain all desired data.

- When all data is read, the cost of an index scan is:

  cost = B tree level I/O + no. of leaf page I/O + cluster count.

- When all data is read but only a leaf scan required, the cost of the index scan is:

  cost = B tree level I/O + no. of leaf page I/O.

- When a row is read, the cost of an index scan is:

  cost = B tree level I/O + one leaf page I/O + one data page I/O.

- When a row is read but only a leaf scan required, the cost of an index scan is:

  cost = B tree level I/O + one leaf page I/O.

- When partial data is read, the cost of an index scan is:

  cost = B tree level I/O + (no. of leaf page x S) + (cluster count x S),

  where S stands for selectivity.

Cost of Sort

Reading data from disk into memory is the only thing that takes more time. The computing cost is proportional to "**c x w x n x** $\mathbf{log_2(n)}$", where c is the number of columns being sorted, w is the bytes of the sorted key, and n is the number of rows being sorted.

Cost of Nested Join

More than two loops are required to access data pages for a nested join. In a nested join, the outer table is different from the inner. Generally, the cost of a nested join is
outer table I/O + inner table I/O x number of rows in outer table

Cost of Merge Join

It is necessary to sort tables before performing a merge join. Suppose two tables have already been sorted for merging with the merge keys. The cost of the merge join is the sum of the I/O for these two tables. If sorting is not performed on the merge keys, the cost of the sorting will still need to be added.

18.4 Statistics

Statistics represent the amount and distribution of data for a table. They provide the information for the cost functions to find the best access plan. However, the statistics will be out of date if data in the table is being inserted, deleted, or updated. The UPDATE STATISTICS command should be used to update statistical values and find real time statistics to enhance the efficiency of a query.

Types of Statistics

DBMaker will collect the following statistics:

FOR A TABLE

- **nPg** – number of data pages
- **nRow** – number of data rows
- **avLen** – average bytes for a row

FOR A COLUMN

- **distVal** – number of distinct values
- **avLen** – average bytes for each column

- **loVal** – the second minimum value for a column
- **hiVal** – the second maximum value for a column

FOR AN INDEX

- **nPg** – number of index pages
- **nLevel** – number of levels in an index tree
- **nLeaf** – number of leaves in an index tree
- **distKey** – number of distinct keys
- **distC1** – number of distinct keys for the first index column
- **distC2** – number of distinct keys for the first two index columns
- **distC3** – number of distinct keys for the first three index columns
- **nPgKey** – number of index pages for each key
- **cCount** – number of clusters counted; the number of data pages accessed through an index

UPDATE STATISTICS Syntax

The *object_list* Clause

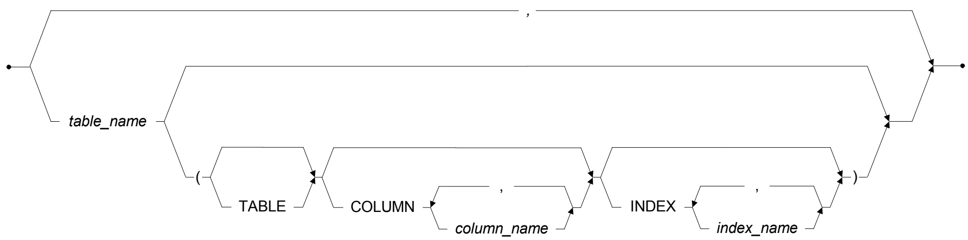

Figure 18-1 Syntax for the UPDATE STATISTICS Statement

- **SAMPLE**-- means the sampling rate expressed as a percentage of the whole, an integer between 1 and 100.

➲ **Example 1:**

To update statistics for all tables including columns, indexes, and system tables with the default sampling rate value of 100:

```
dmSQL> UPDATE STATISTICS;
```

➲ **Example 2:**

To update statistics for all tables including columns, indexes, and system tables with the default sampling rate value of 30:

```
dmSQL> UPDATE STATISTICS sample=30;
```

➲ **Example 3:**

To update statistics for the **jeff.emp** table:

```
dmSQL> UPDATE STATISTICS jeff.emp;
```

➲ **Example 4:**

To update statistics for columns name, age and index **idx1** for the **jeff.emp** table:

```
dmSQL> UPDATE STATISTICS jeff.emp (TABLE COLUMN(name, age) INDEX(idx1));
```

➲ **Example 5:**

To update statistics for the **jeff.emp**, and **jeff.dept** tables:

```
dmSQL> update statistics jeff.emp, jeff.dept;
```

Alternatively, DBMaker provides a date setup to update statistics automatically. Refer to Chapter 6, *Managing Schema Objects* for more information about updating statistics automatically.

Auto Update Statistics Daemon

DBMaker also provides a daemon to automatically update statistics for the entire database. Not all statistics on all tables are regenerated, but rather an optimum sample rate based on how recently statistics were updated for a table and how much the table has been changed since the last time its statistics were updated. Statistics are updated daily at 3:00 AM according to the system clock.

Setting the configuration parameter **DB_StSvr** equal to 1 in the **dmconfig.ini** file activates the Auto Update Statistics Daemon. The keyword must be set before database startup.

Load and Unload Statistics

Users can use the UNLOAD STATISTICS command to dump the statistics values into an external text file. Users can also use the LOAD STATISTICS command to copy statistics values to a database from an external text file.

➲ **Example 1**

To use UNLOAD STATISTICS:

```
dmSQL> unload statistics to file1;//dump statistics from database to file1
```

➲ **Example 2**

To use LOAD STATISTICS:

```
dmSQL> load statistics from file1;//read statistics from external text file
```

An experienced user can enhance the efficiency of query by means of modifying files with statistics, and input it into database.

➲ **Example 3**

An external text file generated from UNLOAD STATISTICS:

```
DBname = TESTDB
```

```
TBowner = jeff
TBname = emp
TBpage = 5
TBrows = 30
Tbavlen = 50

COname = age
COtype = INTEGER
COdist = 12
COavlen = 4
COlow = 25
COhigh = 42

IXname = idxage
IXpages = 5
IXlevel = 2
IXleaf = 3
IXdist = 12
IXdistC1 = 12
IXdistC2 = 12
IXdistC3 = 12
IXpgkey = 8
IXcount = 7
```

18.5 Accelerating Execution of Query

The execution of a query can be accelerated by making the following modifications:

- Read fewer data rows
- Avoid sorting or sort on fewer data rows and columns
- Read data sequentially

Data Model

The definition of a data model includes all tables, views, and indices on the database, especially the existence of indices. It describes whether an index should be used in the conditions such like join, sort, and views.

Query Plan

Use the "Set Dump Plan ON" command to check the query execution plan for DBMaker.

Characteristics of an execution plan:

- **Index**-- checks the output data to see whether an index has been used or not and if so, how to use it.
- **Filter**-- checks the predicate factors to see how much data the predicate can filter.
- **Query—**checks the query after completion to see whether the access plan is the best.

➲ **Example**

Use the following to see the execution plan:

```
dmSQL> set dump plan on;
```

Index Check

Check whether proper indexes on query columns exist. Use methods mentioned in the following sections to improve query efficiency.

Filter Columns

It only takes a small part of source information to make an efficient query. Users can use the WHERE predicate in a SELECT command to control the amount of output information, known as a data filter.

Following are some methods for using the advanced WHERE predicate:

AVOIDING CORRELATED SUB-QUERIES

Correlated sub-queries occur when duplicate columns appear in the main and the sub-query of a WHERE predicate. A different result is returned for a duplicate sub-query for each data row contained in the main query. If the data for columns in each row is different from the one in the previous row in a sub-query, then it is the equivalent to executing a new query for each row gained from the main query.

If you have found a time-consuming sub-query, first check whether it is a correlated sub-query. If so, rewrite the query to avoid the condition. If it is not easy to rewrite the query, try another method to reduce the number of data rows.

AVOIDING DIFFICULT REGULAR EXPRESSIONS

The keyword LIKE provides a comparison of wild cards, known as a regular expression. When a wild card is used at the beginning of an expression, a database server will check each row because it cannot use an index filter. This will instruct DBMaker to sequentially access and check every row in a table.

➲ **Example:**

Using the LIKE keyword with the "*" wildcard:

```
dmSQL> SELECT * FROM emp WHERE name LIKE '*st';
```

Query Results

When you understand what a query really does, you can find another equivalent query to get the same result. We give some suggestions for users to rewrite queries that are more efficient.

- Rewrite joins using views
- Avoid or reduce sorting
- Avoid accessing a large table sequentially
- Use unions to avoid sequential access

Temporary Tables

It is useful to create a temporary, ordered table to accelerate a query. It also can help to avoid sorting operations on multiple columns to simplify the operation of the optimizer.

- Use a temporary table to avoid sorting on multiple columns
- Replace sorting on non-sequential access

18.6 Syntax-Based Query Optimizer

Optimizer chooses a query execution plan with the cost function and statistics automatically. In some special cases when data distribution is skewed, the optimizer may choose a poor query execution plan. To solve the problem, DBMaker supports an optimizer mechanism called the *syntax based query optimizer*.

You can now manually specify the type of scan to use in a query, and the indexes to be used in an index scan. In addition, the DBMaker query optimizer automatically determines the most efficient type of scan to use, even if you have not recently updated the database statistics. There are five different cases where the type of index to be used.

Forced Index Scans

General syntax used to force an index scan:

```
tablename (INDEX [=] idxname [ASC|DESC])
```

➲ **Example 1**

To force a table scan specify the value 0:

```
SELECT * FROM t1 (INDEX=0)
```

➲ **Examples 2:**

To force an index scan on a primary key specify the value 1:

```
SELECT * FROM t1 (INDEX=1)
```

➲ **Example 3**

To force an index scan on the index **idx1**:

```
SELECT * FROM t1 (INDEX idx1)
```

➲ **Example 4**

Allows the query optimizer to decide what type of scan to use on table **t1**, but forces an index scan on the **idx1** index for table **t2**:

```
SELECT * FROM t1, t2 (INDEX idx1)
```

Forced Index Scan and "Alias"

General syntax used to force an index scan and provide an alias for the table:

```
tablename (INDEX [=] idxname) aliasname
```

➲ **Example**

To force an index scan on the **idx1**index, and provides an alias for the table:

```
SELECT * FROM t1 (INDEX idx1) a, t1 b WHERE a.c1 = b.c1
```

Forced Index Scan and "Synonym"

General syntax used to force an index scan using a synonym:

```
synonymname (INDEX [=] idxname)
```

➲ **Example**

To force an index scan on the **idx1** index using synonym **s1**:

```
SELECT * FROM s1 (INDEX idx1)
```

Forced Index Scan and "View"

General syntax used to force an index scan when creating a view:

```
viewname (INDEX [=] idxname)
```

➲ **Example 1**

To force an index scan on the **idx1** index when creating view **v1**:

```
CREATE VIEW v1 as SELECT * FROM t1 (INDEX idx1)
```

You cannot force an index when selecting a view.

➲ **Example 2**

A wrong usage that will return errors:

```
SELECT * FROM v1 (INDEX idx1)
```

Forced Text Index Scans

General syntax used to force a text index scan:

```
tablename (TEXT INDEX [=] idxname)
```

➲ **Example**

To force a text index scan on the **tidx1**index:

```
SELECT * FROM t1 (TEXT INDEX tidx1)
```

18.7 How to Read a Dump Plan

The first step to check a slow query is to read the execution plan. DBMaker supports a dump and read execution plan function.

There are three dmSQL commands for dump plans:

```
dmSQL> Set dump plan on:
```

Turns on the dump plan option. The latter queries will dump plan and then execute commands.

```
dmSQL> Set dump plan off:
```

Turns off the dump plan option. The latter queries will only execute commands but not dump. This is the default option.

```
dmSQL> Set dump plan only:
```

Turns on the dump plan only, but does not execute commands.

At first glance, it appears that a dump plan is composed of several blocks called ON. Query optimizer divides a query into several ON blocks, and each block is a logical optimization unit. The optimizer will then optimize every ON block. Simple and joined queries usually have only one ON block, but a complex query like a sub query may generate more than one ON block, including a main block and sub blocks.

Optimizer finds the best execution method based on cost for every ON block. It will divide an ON block into several PL blocks, where each PL block represents an operation, such as scan, join etc.

Be familiar with the terms introduced in previous sections of this chapter:

- table scan
- index scan
- nested join
- merge join
- factor

Table Scan

➲ **Example:**

To set the dump plan for a table scan of **t1**:

```
dmSQL> set dump plan on;
dmSQL> select * from t1 where c1>1;
```

Result, the dump plan for the table scan of **t1**:

```
----- begin dump plan -----

{ON Block 0}
ON Type     : SCAN

[PL Block 0]
Method      : Scan
Table Name : t1
Type        : Table Scan
Order       : <none>
Factors     : (1) t1.c1 > 1
I/O Cost    : 101.0
CPU Cost    : 25.3
Sub Cost    : 0.0
Result Rows: 330.0
```

```
----- end dump plan -----
```

The first two lines give the information for an ON block:

{**ON Block 0**} -- an ON block with a block ID of 0

ON Type: SCAN -- ON block type is a scan.

The ON block contains one PL block:

[**PL Block 0**] -- A PL block with a block ID of 0.

Method: Scan -- This PL block will perform a scan operation.

Table Name: t1 -- Scan on table t1 defined.

Type: Table Scan -- Scan type is a table scan.

Order: <none> -- Scan order, there is no use for a table scan.

Factors: (1) **t1.c1 > 1** -- This scan will use the filter t1.c1 > 1.

I/O Cost: 101.0 -- Estimated I/O cost is 101.0 pages.

CPU Cost: 25.3 -- Estimated CPU cost is 25.3 pages.

Sub Cost: 0.0 -- Estimated sum of costs for the PL block's sub-block. In this example, there is no PL sub-block.

Result Rows: 330.0 - Estimated result rows after the scan and filter.

Index Scan

➲ **Example:**

To set the dump plan for **c1** and **c2** from table **t2** using WHERE:

```
dmSQL> set dump plan on;
dmSQL> select c1,c2 from t2 where c1>1 and c2=2;
```

Result, the dump plan for **c1** and **c2** from table **t2** using WHERE:

```
----- begin dump plan -----

{ON Block 0}
ON Type     : SCAN
```

```
[PL Block 0]
Method     : Scan
Table Name : t2
Scan Type  : Index Scan on idx21(c2, c1)
Order      : ASC
Index EQFA#: 1
Index FA#  : 2
Index FACOL: 1, 2
Index Cost : 2
Factors    : (1) t2.c2 = 2
           : (2) t2.c1 > 1
I/O Cost   : 2.0
CPU Cost   : 0.6
Sub Cost   : 0.0
Result Rows: 13.0

----- end dump plan -----
```

The first two lines give the information for an ON block:

{**ON Block 0**} – An ON block with a block ID of 0.

ON Type: SCAN -- An ON block type of scan. The ON block also contains one PL block: [**PL Block 0**] -- A PL block with a block ID of 0.

Method: Scan -- The block executes a scan.

Table Name : t2 -- Scan of table t2.

Scan Type: Index Scan on idx21(**c2, c1**) -- Scan type is an index, applying an index to idx12 using the index columns c2 and c1.

Order: ASC -- Ascending index scan order.

Index EQFA#: 1 -- Equal factor number applied in the index scan, in this example using t2.c2 = 2.

Index FA#: 2 -- Factor number that applied in the index scan, in this example using t2.c2 = 2 and t2.c1 > 1.

Index FACOL: 1, 2 –– Factor ID mapping from index columns. In this example, it maps the first index column, c2, to factor, (1) t2.c2=2, and the second index column, c1, maps to factor, (2) t2.c1 > 1.

Index Cost: 2 -- Estimated index page cost of 2.

Factors: (1) t2.c2 = 2

(2) t2.c1 > 1 - Applies filters t2.c2 = 2 and t2.c1 > 1.

I/O Cost: 2.0 -- Estimated I/O cost of 2.0 pages.

CPU Cost: 0.6 -- Estimated CPU cost of 0.6.

Sub Cost: 0.0 -- Estimated sum of costs for the PL block's sub-block.

Result Rows: 13.0 -- Estimated result of rows after the scan and filter.

Equal Join

➲ **Example:**

To set the dump plan from **t1** and **t2** using WHERE:

```
dmSQL> set dump plan on;
dmSQL> select * from t1, t2 where t1.c2=t2.c2;
```

Result, the dump plan from **t1** and **t2** using WHERE:

```
----- begin dump plan -----

{ON Block 0}
ON Type     : JOIN

[PL Block 0]
Method      : Join
Type        : Merge Join
Factors     : (1) t1.c2 = t2.c2
I/O Cost    : 8.5
CPU Cost    : 573.8
Sub Cost    : 231.6
Result Rows: 500.0
Sub Block 1: [PL Block 1]
```

```
Sub Block 2: [PL Block 2]

[PL Block 1]
Method     : Sort
I/O Cost   : 4.2
CPU Cost   : 274.4
Sub Cost   : 120.0
Result Rows: 1000.0
SUB Block  : [PL Block 3]

[PL Block 3]
Method     : Scan
Table Name : t2
Type       : Table Scan
Order      : <none>
Factors    : <none>
I/O Cost   : 101.0
CPU Cost   : 25.3
Sub Cost   : 0.0
Result Rows: 1000.0

[PL Block 2]
Method     : Sort
I/O Cost   : 4.2
CPU Cost   : 274.4
Sub Cost   : 120.0
Result Rows: 1000.0
SUB Block  : [PL Block 4]

[PL Block 4]
Method     : Scan
Table Name : t1
Type       : Table Scan
Order      : <none>
Factors    : <none>
I/O Cost   : 101.0
CPU Cost   : 25.3
Sub Cost   : 0.0
Result Rows: 1000.0
```

```
----- end dump plan -----
```

In this example, there is more than one PL block. The PL block relationship is combined using the sub-block information.

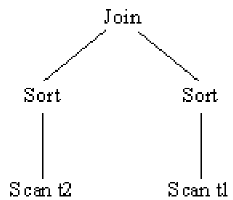

A simple tree representing blocks

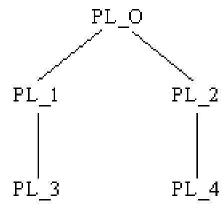

Replace each node with the names

Descriptions of Join blocks:

[**PL Block 0**] -- A PL block with a block ID of 0.

Method: Join -- The block is a Join.

Type: Merge Join -- Join type is merge.

Factors: (**1**) **t1.c2 = t2.c2** -- Apply Join filter t1.c2 = t2.c2 using a Join block.

I/O Cost: 8.5-- Estimated I/O cost is 8.5 pages.

CPU Cost: 573.8 -- Estimated I/O cost is 573.8 pages.

Sub Cost: 231.6 -- Estimated sum of costs for the PL block's sub-block.

Result Rows: 500.0 -- Estimated result rows after the Join block.

Sub Block 1: [**PL Block 1**] – The block's first child links to [PL Block 1].

Sub Block 2: [**PL Block 2**] -- The block's second child links to [PL Block 2].

Description for a sort block:

[**PL Block 1**] -- A PL block with a block id of 1.

Method: Sort -- A sort block.

I/O Cost: 4.2 -- Estimated I/O cost is 4.2 units.

CPU Cost: 274.4 -- Estimated CPU cost is 274.4 units.

Sub Cost: 120.0 -- Estimated sum of costs for the PL block's sub-block.

Result Rows: 1000.0 -- Estimated result rows after the sort block.

SUB Block: [**PL Block 3**] -- This block's child block link to [PL Block 3].

Listed above are the most common dump plan cases that users will encounter. Many changes will be evident in dump plans, but they all consist of the same elements, I/O cost, CPU cost, and Result Rows. If a dump plan is too complex, use the syntax-based optimizer discussed previously to try other methods.

A Keywords in dmconfig.ini

A.1 General Concept

When the DBMaker database engine is started or when a user wants to connect to a database server, DBMaker must initialize several parameters to configure itself. These parameters are read from an ASCII text configuration file named **dmconfig.ini**. This text file contains the keywords and corresponding values that will be used by for configuration. Since this file is in ASCII format, a DBA can edit it with a text editor to change the parameters as required.

In most cases, the keywords are required when a database starts. Keywords changed after the database has been started will not take effect until the database has been shut down and restarted.

However, some of the keywords are only required when users connect to database servers. Users can change these keywords after the server has started, but before the connect command has been issued, if they want the new values to be used during the current session.

The configuration parameters play an important role in the performance of DBMaker. To ensure DBMaker will run smoothly, be aware of the effects of each configuration parameter and estimate the best values to use. It is recommended that the database administrator back up the **dmconfig.ini** file as well as the database files on a regular basis.

A.2 dmconfig.ini File Format

The **dmconfig.ini** file is an ASCII text file. It can be edited with any text editor capable of opening and saving ASCII text files. The **dmconfig.ini** file consists of many sections, each section is made up of the configuration information used to start a particular database. Each section begins with the section name, followed by a list of keywords and their values.

➲ **Example**

The format of a dmconfig.ini file:

```
[section_name_1]
keyword1 = value1          ; here is a comment
keyword2 = value2
         .
         .
[section_name_2]
keyword3 = value3 value4   ; spaces or commas may be used
keyword4 = value5          ; as delimiters
         .
         .
```

Section Names

The name of each section corresponds to the name of the database that will use the configuration options found in that section when it starts up. The section name begins with a left square bracket ([) followed by the name of the database, and ends with a right square bracket (]). The brackets are required to enclose the section name, and the left bracket must be the first character on that line.

Keywords

Following each section name is a list of keywords and their values. These values will be used by the database that corresponds to the section heading for configuration when it starts. The statement keyword=value assigns the specified value to a keyword. The value of a keyword can be an integer or a string, depending on the keyword itself.

Comments

Any string or symbol that is written after the semi-colon (;) will be considered a comment and ignored by DBMaker.

➲ **Example,**

dmconfig.ini file using (;) to include comments:

```
[SDB]
DB_DbFil=SDB.DB
DB_JnFil=SDE.JNL
DB_SMode=1              ;normal start mode
DB_UMode=0              ;single-user
DB_BMode=1
DB_BkSvr=1
DB_BkTim=96/03/19 00:00:00
DB_BkItv=7-00:00:00
DB_NBufs=100
DB_NJnlB=200
DB_MaxCo=100
DB_JnlSZ=20000
file1=SDE.FIL 40

[EMP]
DB_DBFIL=EMP.DB
DB_JNFIL=EMP.JNL
DB_SMode=1              ;normal start mode
DB_UMode=1              ;multi-user
DB_NBufs=100
DB_NJnlB=400
DB_MaxCo=100
```

```
DB_DbfSz=50
DB_JnlSZ=20000
file1=EMP.FL1 100
file2=EMP.FL2 200
```

In the example, the **dmconfig.ini** contains two sections one for the **SDB** database and the other for the **EMP** database.

A.3 Search Path for dmconfig.ini

Another issue for **dmconfig.ini** is where to locate this file. For UNIX systems, DBMaker will search in three locations for the **dmconfig.ini** file.

➲ **The locations and the search order are:**

1. The current directory.
2. The directory specified in the environment variable DBMAKER.
3. UNIX systems: the subdirectory data in the home directory for the user dbmaker (~dbmaker/data).

In Microsoft Windows systems, the **dmconfig.ini** file must be placed in the installation directory ;usually the **WINDOWS** directory, unless you changed the default when installing Windows.

When starting a database, DBMaker will scan the above UNIX directories or the **WINDOWS** directory in the order listed to locate a **dmconfig.ini** file with a section name that corresponds to the database. If a **dmconfig.ini** file is found and the section name exists, the keywords defined in the section will be used. If the section name cannot be found in that file, DBMaker will continue searching for a **dmconfig.ini** file sequentially in the directories until the section name is found.

A.4 Default Values for Keywords

Whenever DBMaker needs a parameter, it will search the corresponding keywords in the proper section of the **dmconfig.ini** file. The pattern matching for section names or keywords in a search are case insensitive, except for user-defined files. If a keyword cannot be found in the **dmconfig.ini** file, a default value will be used. Most of the keywords have their own default values. Refer to the last section in this appendix for more information.

A.5 Creating dmconfig.ini

Normally a database administrator creates the corresponding section in the **dmconfig.ini** file with a text editor before a database is created and the parameters take effect while creating that database. However, if DBMaker cannot find the section in **dmconfig.ini** while creating a database, it will automatically create a section in the first **dmconfig.ini** file found or in a new **dmconfig.ini** file, if not found. Therefore, a section for each database should always be found at startup time, if not found, DBMaker will return an error.

A.6 Keyword Reference

DB_AtCmt=<value>

This keyword specifies whether auto-commit mode is on or off. Setting this value to 1 turns auto-commit mode on and setting it to 0 turns auto-commit mode off. When auto-commit mode is on, DBMaker will automatically issue a COMMIT TRANSACTION after each SQL command is successfully executed. This keyword is set from the client side.

default value: 1

valid range: 0, 1

see also: **DB_LTimO**

where to use: **client side**

DB_AtrMd=<value>

This keyword specifies whether the database is a source database of asynchronous table replication. Setting this value to 1 turns on the logging of base table operations and also enables the Distributor Daemon. Therefore, it can be a source database of a replication.

default value: 0

valid range: 0, 1

see also: **DB_EtrPt**, **RP_LgDir**

where to use: **server side**

DB_BbFil=<string>

This keyword specifies the name of the system BLOB file. It will expand as necessary as more BLOB data is inserted into this file.

default value: database name with the file extension ***.SBB***. For example, ***db.SBB***.

valid range: string with length < 256

see also: **DB_DbDir**, **DB_DbFil**, **DB_UsrBb**, **DB_UsrDb**

where to use: **server side**

DB_BfrSz=<value>

This keyword specifies the size of each BLOB frame in kilobytes. This keyword is used when the database is created.

default value: 16 (KB) for UNIX, Windows 98, Windows NT

valid range: 8~256 for UNIX, Windows 98,Windows NT

see also: **DB_BbFil**

where to use: **server side (only for creating a database)**

DB_BkDir=<string>

This keyword specifies the directories where the backup server puts the database backup files. These directories must already exist and can be different from **DB_DbDir**.

valid range: string with length < 256

see also: **DB_BkSvr**, **DB_BMode**

where to use: **server side**

➲ **Example**

< BKDIR n >: the n's backup path

< SIZE n >: the size of the n's backup path (4k per unit)

DB_BKDIR = <BKDIR 1> <SIZE 1> < BKDIR 2> <SIZE 2> < BKDIR 3> <SIZE 3>…

```
[MYDB]

…..
DB_BKDIR = /home/usr/dbmaker/bk 5000 /home2/backup 1000
…..
```

When /home/usr/dbmaker/bk is full, it will backup files to /home2/backup.

DB_BkFoM=<value>

This keyword specified the file object (FO) backup mode. File objects will only be backed up during a full backup of the database. **DB_BkFoM** has three possible values; 0, 1, and 2. Setting **DB_BkFoM** equal to zero disables the FO backup feature; file objects will not be backed up during a full backup. Setting **DB_BkFoM** equal to one enables system file objects to be backed up during a full backup. Setting **DB_BkFoM** equal to two enables both system file objects and user file objects to be backed up.

default value: 0

valid range: 0: File objects are not backed up

1: System file objects are backed up

2: System and user file objects are backed up.

see also: **DB_BkSvr, DB_FBKTm, DB_FBKTV, DB_BkDir.**

where to use: **server side**

DB_BkFrm=<value>

The keyword allows you to specify the format Backup Server used to name incremental backup journal files. The backup filename format may include both text constants and format sequences (escape sequences), that represent special character strings.

You can use the format sequences to represent the year, month, or date the backup was performed, the name of the database, or the backup identification number. You may combine format sequences and text constants in any way, provided the result is a valid filename supported by the operating system. The format sequences have three parts: the escape character, the length value, and the format character. The valid format sequences are:

%[***n***]**Y**—The backup year of the journal file was.

%[***n***]**M**—The backup month of the journal file was.

%[***n***]**D**—The backup day of the journal file was.

%[***n***]**B**—The backup identification number.

%[***n***]**N**—The name of the database the journal file belongs to.

➲ **Example**

DB_BkFrm = %N.%B

If the database is test1, the incremental backup files will be named: test1.1, test1.2…

default value: **%4N%4B.jnl**

see also: **DB_BkSvr, DB_BkTim, DB_BkItv**

where to use: **server side**

DB_BkFul=<value>

This keyword specifies the percentage that all journal files must be filled to before the backup server is triggered to do an incremental backup. Setting this value to 0 will trigger the backup server whenever a journal file is full. Setting this value between 50-100 will trigger the backup server whenever the total space used in all of the journal files exceeds the specified percentage. For example, if there are two journal files of 500 journal blocks each and **DB_BKFUL** is set to 80, then after every $500 \times 2 \times 0.8 = 800$ blocks are used, the backup server will automatically do an incremental backup.

default value: 0

valid range: 0, 50 ~ 100

see also: **DB_BkSvr**, **DB_BkTim**, **DB_BkItv**

where to use: **server side**

DB_BkItv=<string>

This keyword specifies the backup time interval. Please refer to **DB_BkTim** described later.

default value: none (no backup schedule if **DB_BkItv** is not set)

see also: **DB_BkSvr**, **DB_BkTim**, **DB_BMode**

where to use: **server side**

DB_BkCmp=<value>

This keyword specifies whether the compact backup mode is used. Not every journal block in a journal file is needed to perform a backup. If this keyword is set to 1 the backup server will only back up those journal blocks that require back up to save disk space. Also see the chapter *Database Backup, Recovery, and Restoration*.

default value: 1

valid range: 0,1

see also: **DB_BkSvr**

where to use: **server side**

DB_BkOdr=<string>

This keyword specifies the directories where the backup server puts the pervious version of full backup files. Also, refer to Chapter 14, "Database Backup, Recovery, and Restoration".

default value: none.

valid range: string with length < 256

see also: **DB_BkSvr**, **DB_BMode**, **DB_FBkTm**, **DB_FBkTv**

where to use: **server side**

➲ **Example**

DB_BKODR = <BKDIR 1> <SIZE 1> < BKDIR 2> <SIZE 2> < BKDIR 3> <SIZE 3>…

< BKDIR n > : the n's backup path

< SIZE n > : the size of the n's backup path (4k per unit)

```
[MYDB]

…..
DB_BKODR = /home/usr/dbmaker/bk 5000 /home2/backup 1000
…..
```

DB_BkSvr=<value>

This keyword specifies whether or not a backup server will be started when a database is started. Setting this value to 1 will start a backup server for that database. Also, refer to Chapter 14, "Database Backup, Recovery, and Restoration".

default value: 0

valid range: 0,1

see also: **DB_BkCmp**, **DB_BkDir**, **DB_BkFul**, **DB_BkTim**, **DB_BkItv**

where to use: **server side**

DB_BkTim=<string>

This keyword along with **DB_BkItv** specifies the schedule of the backup server. **DB_BkTim** specifies the first time a backup server will perform an incremental backup. Incremental backup will then be performed after every time interval specified in **DB_BkItv**.

➲ **Example**

```
DB_BkTim = 96/05/01 00:00:00    ;backup begins from May 1, 1996.
DB_BkItv = 1-12:30:00           ;interval is every one day, 12 hours
                                 and 30 minutes.
```

The keywords **DB_BkTim** and **DB_BkItv** are meaningful only when the backup server is started. Also, see the chapter on "Database Backup, Recovery, and Restoration".

default value: none (no backup schedule if **DB_BkTim** is not set)

see also: **DB_BkItv**, **DB_BkSvr**, **DB_BMode**

where to use: **server side**

DB_BMode=<value>

This keyword specifies the backup mode of a database. Setting the value to 0 enables NON-BACKUP mode, 1 enables BACKUP-DATA mode, and 2 enables BACKUP-DATA-AND-BLOB mode. Also see the chapter *Database Backup, Recovery, and Restoration*.

default value: 0

valid range: 0 ~ 2

see also: DB_BkSvr

where to use: server side

DB_Brows=<value>

This keyword specifies the lock behavior of a select statement. Setting the value to 0 denotes DBMaker will take S lock on the result set of select statement, and 1 denotes DBMaker will not lock the result set of select statement. This value is required while connecting to a database.

default value: 1

valid range: 0,1

where to use: client side

DB_CBMod=<value>

This keyword specifies the behavior of the cursor after the end of a transaction. A value of 1 indicates all still open cursors will be closed after any transaction is committed. A value of 2 or 3 indicates all still open cursors will be kept open after a transaction is committed. 2 indicates that all locks would be released after the end of the transaction. A 3 indicates that all locks would be reserved but all exclusive ones would become shared. In those cases (1, 2 or 3), the cursor will be closed if any transaction is aborted.

default value: 2

valid range: 1 ~ 3

where to use: **client side**

DB_CmChe=<value>

This keyword specifies whether client command cache is turned on or off. Setting this value to 1 turns on client command cache and setting it to 0 turns it off. When client command cache is turned on, the related information of previous executed SQL command can be kept in cache and reused in the following command execution if the SQL command is the same. By this way, better performance is gained. This keyword is set from the client side.

default value: 1

valid range: **0**, 1

where to use: **client side**

DB_CTimO=<value>

This keyword specifies the connection time-out value in seconds when a client is trying to connect to the database server. If a database has not been started or the server IP address is wrong, users may be forced to wait a long time until the connection times out. Users can set the value of this keyword to shorten the waiting time. This parameter is set from the client side.

default value: 5 (seconds)

valid range: 5 ~ 1:00:00 (1 hour)

where to use: **client side**

DB_DaiFm=<value>

This keyword specifies the date input format for SQL statements. Refer to "*Appendix B"* in the "*ODBC Programming Guide"* for more information.

default value: none (accept all date input formats)

valid range: **mm/dd/yy**

mm-dd-yy

dd/mon/yy

dd-mon-yy

mm/dd/yyyy

mm-dd-yyyy

yyyy/mm/dd

yyyy-mm-dd

dd/mon/yyyy

dd-mon-yyyy

dd.mm.yyyy

see also: DB_DaoFm

where to use: client or server side (client has higher priority)

DB_DaoFm=<value>

This keyword specifies the date output format in SQL statements. Refer to the "ODBC Programmer's Guide", Appendix B for more information.

default value: **yyyy-mm-dd**

valid range: same as valid range of DB_DaiFm

see also: DB_DaiFm

where to use: **client or server side (client has higher priority)**

DB_DbDir=<string>

This keyword specifies the directory that the database files reside in. The directory string can be a relative or a full path name.

There are seven types of files in DBMaker;

- system database file defined by **DB_DbFil**
- default user data file defined by **DB_UsrDb**
- system journal file defined by **DB_JnFil**
- system BLOB file defined by **DB_BbFil**
- default user BLOB file defined by **DB_UsrBb**
- system temporary files defined by **DB_TpFil**
- user defined files

When defining full path names for these keywords, relative path names or simple file names can be used. If a path name is used in defining the keywords, DBMaker will use that name to reference the defined file. If a simple file name is used, DBMaker will search for the **DB_DbDir** keyword. If this keyword is found, DBMaker will prefix the simple file name with the string specified in **DB_DbDir** to reference the file. If it is not found, DBMaker will use the file name and assume it is located in the current directory.

➲ **Example 1**

```
[DB1]
DB_DbDir = /disk1/db
DB_DbFil = mydb1
DB_JnFil = /disk2/usr/DB1.JNL
```

Using the physical file names:

```
DB_DbFil -- /disk1/db/mydb1
DB_JnFil -- /disk2/usr/DB1.JNL
DB_BbFil -- /disk1/db/DB1.BB (using default file name)
```

➲ **Example 2**

```
[DB2]
DB_DbFil = mydb2
DB_JnFil = /disk2/usr/DB2.JNL
```

Using the physical file names:

```
DB_DbFil -- mydb2 ( in current directory )
DB_JnFil -- /disk2/usr/DB2.JNL
DB_BbFil -- DB2.BB ( in current directory )
```

default value: (current directory)

valid range: string with length < 256

see also: **DB_DbFil, DB_JnFil, DB_BbFil, DB_TpFil, DB_UsrDb, DB_UsrBb**

where to use: **server side**

DB_DbFil=<string>

This keyword specifies the physical name of the system database file used by the operating system.

default value: database name with the file extension **.SDB**. For example, **db.SDB**.

valid range: string with length < 256

see also: **DB_BbFil, DB_DbDir, DB_UsrBb, DB_UsrDb**

where to use: **server side**

DB_DifCo=<value>

This keyword specifies the connection behavior of an application connecting to the same database more than once. Setting the value to 1 instructs DBMaker to treat each duplicate connecting action as a separate connection. Setting this value to 0 instructs DBMaker to merge all duplicate connecting actions as one connection. This setting is used while connecting to a database.

default value: 1

valid range: 0,1

where to use: **client side**

DB_DtClt=<value>

This keyword specifies the time interval DBMaker uses to automatically detect the existence of a client. Setting this keyword to 0 disables this feature. Sometimes when a client machine is suddenly powered off or the network is not working properly, the server cannot release the resources allocated to the client. Turning this keyword on will solve the problem because the server will auto-detect the client. If the server finds a dead client, it will release all of its resources.

default value: 300 (seconds)

valid range: 0, 5 ~ 1:00:00 (1 hour)

see also: **DB_ITimO**

where to use: **client side**

DB_ERMRv=<string>

This keyword specifies recipients for e-mail that is generated when the database experiences an error. Up to eight recipient e-mail addresses may be specified; a comma must separate each address string. If no recipient is specified, the error report system is disabled and no e-mail will be generated.

default value: null

see also: **DB_ ERMSv**

where to use: **server side**

DB_ERMSv=<string>

This keyword specifies the SMTP server to use for relaying e-mail messages. Only one SMTP server may be specified. If no SMTP server is specified but recipients are specified, then DBMaker will assign the 'localhost' value to this keyword.

default value: localhost

see also: **DB_ERMRv**

where to use: **server side**

DB_EtrPt=<value>

This keyword is an integer, which specifies the TCP/IP port number that the Subscriber Daemon for the database server is attached to. This is used for express asynchronous table replication. In the source database, this keyword instructs how to connect to the subscriber of the destination database. In the destination database, this keyword will start the Subscriber Daemon.

default value: none

valid range: 1024~65535

see also: **DB_AtrMd**, **RP_LgDir**

where to use: **server side (both for the source and destination of databases)**

DB_ExtNp=<value>

This keyword specifies a size for DBMaker to extend autoextend tablespace. When an autoextend tablespace is exhausted, DBMaker will automatically extend. The value specifies how many pages/frames to extend.

default value: 20 (pages/frames)

valid range: 1 or higher

where to use: **server side**

DB_FBkTm=<string>

This keyword combined with **DB_FBkTv,** specifies the schedule of the Backup Server to perform an on-line full backup. **DB_FBkTm** specifies the first time the Backup Server will perform a full backup. On-line full backup will be performed after every time interval specified in **DB_FBkTv**.

➲ **Example**

```
DB_FBKTm = 96/05/01 00:00:00    ;begins from May 1, 1996.
DB_FBkTv = 1-12:30:00           ;interval is every one day, 12 hours
                                 and 30 minutes.
```

The keywords **DB_FBkTm** and **DB_FBkTv** are only used with the Backup Server.

default value: none

see also: **DB_FBkTv**, **DB_BkSvr**, **DB_BkOdr**

where to use: **server side**

DB_FBkTv=<string>

This keyword specifies the full backup time interval. Refer to **DB_FBkTm** for more information.

default value: none (no full backup schedule if **DB_FBkTv** is not set)

see also: **DB_BkSvr**, **DB_FBkTm**, **DB_BkOdr**

where to use: **server side**

DB_FoDir=<string>

This keyword specifies the path for system file objects or system file object subdirectories (depending on the setting of **DB_FoSub**). You must specify a full path name for **DB_FoDir**.

➲ **Example 1**

```
DB_FoDir = /usr/DBMaker/fileobj          ;for UNIX
```

- **Example 2 :**

```
DB_FoDir = c:\dbmaker\fileobj              ;for DOS
```

- **Example 3**

```
DB_FoDir = \\NTMachine\dbmaker\fileobj     ;for Microsoft UNC name
```

You can insert data into a new system file object only when you have set **DB_FoDir**. This keyword is used when a database is started.

default value: null string (the system file object cannot be inserted)

valid range: string with length < 256

see also: **DB_UsrFo, DB_FoSub**

where to use: **server side**

DB_ForcS=<value>

This keyword instructs DBMaker to force start a database even if an error occurs. Setting the value to 1 enables forced startup.

default value: 0

valid range: 0,1

see also: **DB_SMode**

where to use: **server side**

DB_ForUX=<value>

This keyword specifies the lock behavior of "select … for update" statement in server site.

DBMaker takes S locks on result sets of "select … for update" statements. Special applications that set this value to 1 instruct DBMaker to take X locks on result sets of "select … for update" statements. This setting is required while starting the database.

default value: 0

valid range: 0,1

where to use: server side

DB_FoSub=<value>

This keyword specifies the maximum number of file objects that may be stored in each system file object subdirectory. Subdirectories are created in the directory specified by **DB_FoDir**. A new file object subdirectory will be automatically created when the number of file objects in a subdirectory exceeds the threshold value.

➲ **Example 1**

To set the number of files per file object subdirectory to 500,

```
DB_FoSub = 500
```

You can insert data into a new system file object only when you have set **DB_FoDir**. This keyword is used when a database is started.

default value: 0 (file objects are stored directly in the directory specified by **DB_FoDir**)

valid range: 100 ~ 10000 or 0 (FO directory has no sub-directories)

see also: DB_FoDir

where to use: server side

DB_FoTyp=<value>

This keyword specifies what ODBC type the FILE type maps. ODBC does not define FILE type that is supported only by DBMaker, development tools, such as Borland Delphi or Microsoft Visual Basic do not know about FILE type. If you want to allow tools to access data of the FILE type, DBMaker should internally map the FILE type to LONG VARBINARY by setting DB_FoTyp to 1. There is no mapping if **DB_FoTyp** is 0.

default value: 1

valid range: 0 (not mapping)

1 (FILE type mapping to LONG VARBINARY)

where to use: client side

DB_GcChk=<value>

This keyword is used to set the minimum number of transaction per second (TPS) that will initiate the group commit transaction protocol.

DBMaker always checks the current number of transactions. The number of transactions per second is equivalent to the number of sync requests per second. DBMaker uses **DB_GCCHK** as the TPS threshold. When the server TPS reaches this threshold, the group-commit protocol is activated. When TPS is below the threshold, the group-commit is turned off. For example, if **DB_GCCHK** = 20, the group-commit protocol will be turned on if there are over 20 transactions per second.

DB_GCCHK allows DBMaker to dynamically switch between activating and deactivating the group-commit protocol. Since transaction activity is not constant, e.g. sometimes very high, sometimes low, DBMaker will switch the protocol to avoid unnecessary waiting.

default value: 20

valid range: >0

see also: **DB_GcWtm, DB_GcMxw**

where to use: server side

DB_GcMxw=<value>

This keyword specifies the maximum number of a transactions waiting to do group-commit allowed. It works in conjunction with **DB_GcWtm**, which affects the maximum waiting time for group-commit.

If the group-commits protocol is activated (refer to **DB_GcChk** for details), DBMaker will check for every sync request. If it is meets one of the following conditions, tube sync process will proceed, otherwise it will wait for another sync request before performing the group-commit.

- Transactions reach the maximum waiting time specified by **DB_GcWtm** keyword;
- The numbers of transactions waiting to do group-commit exceeds the value specified by **DB_GcMxw** keyword.

➲ **Example**

The max waiting time is 30 milliseconds and the number of transactions waiting is 5. DBMaker will perform the sync operation when there is at least one transaction waiting over 30 milliseconds or when there are 5 transactions waiting.

default value: 5 (number of transactions)

valid range: > 0

see also: **DB_GcChk, DB_GcWtm**

where to use: server side

DB_GcWtm=<value>

This keyword specifies the maximum waiting time when any transaction waits for group-commit. The longer waiting time may reduce the respond time of the transaction, but may increase the whole throughput for group-commit.

This keyword works with DB_GcMxw. Refer to DB_GcMxw for detail.

default value: 30 (milliseconds)

valid range: > 10

see also: **DB_GcChk, DB_GcMxw**

where to use: server side

DB_IFMem=<value>

This keyword specifies the maximum amount of memory in MB that DBMaker will allocate for use by the IVF text index search routine. DBMaker dynamically manages the amount of memory by allocating half the memory the operating system determines is available. If DBMaker cannot determine the amount of memory available, or detects that more than 128 MB are available, it will set the maximum amout of memory to allocate at 64 MB. Manually adding this keyword and entering a value ansures that DBMaker will never use more than the specified amount. This is recommended if your operating system does not provide memory usage information, or if you want to allocate a larger maximum cache size for improved IVF text index query performance.

➲ **Example 1**

To specify a maximum amount of 256 MB to use for the IVF text index buffer:

```
DB_IFMem = 256
```

default value: No keyword specified – DBMaker automatically configures the buffer

valid range: 64 ~ (maximum allowed by operating system)

see also:

where to use: server side

DB_IDCap=<value>

This keyword specifies the case sensitivity of all identifiers in a database. Setting the value to 0 indicates that all identifiers are case sensitive. Setting the value to 1, indicates that all identifiers in a database are case insensitive; under this mode, all identifiers are converted to uppercase when defined. This keyword can only be set before database creation, which means changing this keyword value for an existing database has no effect.

default value: 1

valid range: 0 (case sensitive)

1 (**case insensitive**)

where to use: **server side**

DB_IOSvr=<value>

This keyword specifies if DBMaker should turn the I/O server and checkpoint daemon on or off. Setting the value to 1 starts the I/O and checkpoint daemons after starting the database server. The keyword is used when the database is started.

default value: 1

valid range: 0, 1

see also: **DB_NBufs**

where to use: **server side**

DB_ITimO=<value>

This keyword specifies the idle timeout interval, specified in seconds. DBMaker will automatically disconnect connections that have no database operations with a higher value than the specified timeout interval. The feature forces idle connections to release all database resources, including buffers, pages, locks, and memory. Set the value to 0 to disable the feature. DBMaker will automatically reset this value to **DB_DtClt** if the value is less than **DB_DtClt**.

default value: **0** (disable)

valid range: **0** ~ **4294967** (seconds)

see also: **DB_DtClt**

where to use: **server side**

DB_JnFil=<string>

This keyword specifies the names of the system journal files. It also specifies the amount of journal files allocated for the database. Up to eight journal files can be specified.

default value: database name with the file extension ***.JNL***. For example, ***DB.JNL***.

valid range: string with length < 256

see also: **DB_JnlSz**, **DB_DbDir**

where to use: **server side**

DB_JnlSz=<value>

This keyword specifies the journal file size using the number of pages (4 KB per page).

default value: 1000 (default journal size is 4 MB)

valid range: 100~524287 and greater than **DB_NJnlB** by at least three

see also: **DB_JnFil**, **DB_DbDir**

where to use: **server side**

DB_LbDir=<string>

This keyword specifies the directory where the user-defined functions (UDF), *dll* in Windows, or *so* in Unix files should be loaded when the database starts.

default value: DBMaker working directory

valid range: string with length < 256

see also: **DB_JnlSz**, **DB_DbDir**

where to use: **server side**

DB_LCode=<value>

This keyword specifies the language code. The language code will affect the result of LIKE operations in a query. A 0 indicates the language is ASCII compatible. A 1 indicates the language code is Chinese BIG5 code compatible. A 2 indicates the language code is SHIFT JIS code compatible. A 3 indicates the language code is GB code compatible. Please refer to the *SQL Reference Manual* for more information. This value is required when the database is started.

default value: **(set during setup)**

valid range: 0 **English (ASCII)**

1 Traditional Chinese (BIG5)

2 Japanese (Shift JIS + Half Corner)

3 Simplified Chinese (GBK)

4 Latin1 code (ISO-8859-1)

5 Latin2 code (ISO-8859-2)

6 Cyrillic code (ISO-8859-5)

7 Greek code (ISO-8859-7)

where to use: **server side**

DB_LetPT=<value>

This keyword specifies the *Lock Escalation Threshold* for escalating a page lock to a table lock. When the number of locks on pages in the same table exceeds the lock escalation threshold, DBMaker will automatically escalate the lock to a table lock.

default value: 50

valid range: 3-32767

see also: **DB_LetRP**

where to use: **server side**

DB_LetRP=<value>

This keyword specifies the *Lock Escalation Threshold* for an escalating rowlock to a page lock. When the number of locks on rows in the same table exceeds the lock escalation threshold, DBMaker will automatically escalate the lock to a page lock.

default value: 15

valid range: 3-32767

see also: DB_LetPT

where to use: server side

DB_LTimO=<value>

This keyword is an integer that specifies the lock time-out value in seconds. When you need to acquire a lock on a database object, such as a table or a tuple, and that object has been already been allocated to another transaction, wait until the object is released.

However, if you prefer not to wait, set this keyword to a desirable value. DBMaker will wait for the object until the waiting time expires, which will return a "lock time-out" error message, or until a lock is acquired on the object before the lock time-out. To disable the time-out, set this keyword to -1. This will cause DBMaker to continue waiting until the lock is released.

Alternatively, set the keyword to 0, indicating that you don't want to wait at all. This keyword is used at connection time rather than database startup time. Each connection may have its own **dmconfig.ini** file, especially in client/server mode, so each user can have a lock time-out value.

default value: 5

valid range: -1~ 65535

where to use: client side

DB_MaxCo=<value>

It denotes the *hard connection number* while creating database or starting database with new journal operation, and denotes the *soft connection number* upon starting a database normally. The *soft connection number* indicates the maximum number of transactions that can be simultaneously active in the database system. It also indicates the maximum number of simultaneous connections that can be established to the database system, since a connection can own at most only one transaction.

The *hard connection number* determines the layout of the journal file regarding the number of connections that it can keep a record of. The hard connection number must be a value between 240 and 1200 and a multiple of 40, so **DB_MaxCo** is rounded up to the closest multiple of 40 (or 240) to determine the hard connection number. For more information regarding the hard and soft connection numbers and their effect on database performance, refer to section 17.5, “Tuning Concurrent Processes”.

default value: 32

valid range: 2~1200

see also: **DB_UMode, DB_SMode**

where to use: **server side**

DB_NBufs=<value>

This keyword specifies the number of data buffers. Each buffer in DBMaker is 4 KB. DBMaker will run more efficiently with more buffers in most cases. This keyword is used when the database is started.

Setting the value to 0 for platforms where DBMaker can detect the physical memory usage allows DBMaker to configure the buffer size automatically. If DBMaker fails to get physical memory information, DBMaker will set the buffer size to 500 pages for Win95/98 or 2000 pages for WinNT/Unix/others.

➲ **Example**

After starting the database, query the SYSINFO table to determine the number of buffers that the database uses. The following command shows that the currently running database uses 500 page buffers:

```
dmSQL> select * from SYSINFO where INFO = 'NUM_PAGE_BUF';

 ID       INFO            VALUE
==== ============== ==================
0107  NUM_PAGE_BUF                 500

1 rows selected
```

For information on determining the optimal value, please refer to section “Tuning Memory Allocation” in Chapter17, “Performance Tuning”.

default value: 0 (auto-configure)

valid range: 0, 15~depends on system

see also: **DB_NJnlB**, **DB_ScaSz**

where to use: **server side**

DB_NetEc=<value>

This keyword specifies if DBMaker should turn network encryption on or off. If network encryption is turned on, all network data between DBMaker server and clients will be encrypted. The encryption technique used by DBMaker is a mix of DES and RSA.

default value: **0** (off)

valid range: **0** (off) , **1** (on)

where to use: **server side**

DB_NJnlB=<value>

This keyword specifies the number of Journal buffers, 4 KB per buffer, in shared memory. This is not the size of the Journal file, but the number of 4KB Journal buffers, which are stored in memory. The keyword is used when the database is started.

default value: 64 (256 KB)

valid range: 16~depends on system

see also: **DB_JnlSz**, **DB_NBufs**, **DB_ScaSz**

where to use: **server side**

DB_Order=<string>

This keyword indicates the name of the order definition file that is put in the **shared/codeorder** subdirectory of DBMaker's installation directory. The order definition file is a pure text file, which would affect the sorting result in DBMaker. This keyword is used when the database is created and then it is rendered useless. Without this keyword, the sorting sequence is in binary sequence.

default value: none

valid range: file name of the user-defined order definition file

see also: **DB_LCode**

where to use: **server side (only for creating a database)**

DB_PasWd=<string>

This keyword specifies the password for the default login user ID. If no default login user ID is specified, the value is ignored. The keyword is used when the database starts or at connection time.

default value: null string

valid range: string with length <= 8 *see also:* DB_UsrId

where to use: client side

DB_PtNum=<value>

This keyword is an integer specifying the TCP/IP port number that the database server is using. It is used at connection time on the client side and at startup time on the server side. The number must match on all client and server machines for a database, or the connection will fail.

default value: none

valid range: 1024~65535

see also: DB_SvAdr

where to use: both client and server sides

DB_ResWd=<value>

Users may need to add (create or import) objects to the database that use DBMaker reserved words as identifiers (see the *SQL Command and Function Reference* for a full list of reserved words). Attempting to add an object that uses a reserved word as an identifier will return an error if DB_ResWd is set equal to 1. If DB_ResWd is set equal to 0, DBMaker will not return an error. The primary purpose of this keyword is to allow objects to be imported that contain reserved words.

default value: 1

valid range: 0, 1

where to use: server side

DB_RmPad=<value>

This keyword specifies whether the space padding for CHAR type data is removed. A value of 0 indicates the space padding for all CHAR type data in a result set is kept. A value of 1 indicates the space padding for all CHAR type data to be removed before copying to a user buffer. It allows a user application to retrieve fixed length CHAR data, excluding the trailing space padding generated in the DBMS during a data insert.

default value: 0

valid range: 0, 1

where to use: **client side**

DB_RTime=<string>

This keyword specifies the target time for a database to be restored from a backup. When performing a database restoration, DBMaker will *roll* forward on the backup files from the earliest time in the to the time specified by **DB_RTime**. If **DB_RTime** is not given, DBMaker will restore the database to the latest time in the backup files, which is the time the backup was performed.

If the **DB_RTime** is later than the backup time, the backup time will be used as the value for **DB_RTime**.

The format for **DB_RTime** is **`yy/mm/dd hh:mm:ss`**.

default value: 0 (70/1/1 00:00:00)

where to use: **server side**

DB_ScaSz=<value>

This keyword specifies the size of the System Control Area (SCA) in kilobytes. If the minimum memory required by DBMaker for the SCA is larger than the value of **DB_ScaSz**, DBMaker will ignore the value and allocate the minimum memory required for the SCA. This keyword is used when the database is started.

default value: 200 (KB) on UNIX, Windows 98, Windows NT

valid range: 1~(system dependent)

see also: **DB_NBufs**, **DB_NJnlB**

where to use: server side

DB_SMode=<value>

This keyword indicates the database startup mode. There are 6 startup modes:

- **normal startup**—Starts a system normally. If the database crashed, DBMaker will perform crash recovery automatically to bring it to a consistent and stable state.
- **startup with new journal**—Starts a system normally, but creates a new journal file with a name given by the value for the ***DB_JNFIL*** keyword. If a file with the same name already exists, it will be overwritten.
- **startup with rollover**—Uses the backup database files, including the journal file, to start the database. DBMaker will rollover the operations to the point in time specified by **DB_RTime**. This mode is used for database restoration. See the chapter on *"Database Recovery, Backup, and Restoration*" for more detailed information about rollover.
- **startup as a primary database**—This mode is used for database replication. Starting a system with this mode makes it a primary database, i.e. a source database. Refer to the chapter on "Database Replication" for more information.
- **startup as a slave database**—This mode is used for database replication. Starting a system with this mode makes it a slave database, i.e. a destination database. For more detailed information refer to the chapter on *"Database Replication*".
- **startup as a read-only database**— Starts up a system normally, but the database is Read-Only or only provides the read privilege. Starting a database with write permissions in Read-Only mode prohibits users from modifying it.

default value: 1

valid range: 1 (normal startup)

2 (startup with new journal)

3 (startup with rollover)

4 (startup as a primary database)

5 (startup as a slave database)

6 (startup as a read-only database)

see also: DB_ForcS

where to use: **server side**

DB_SQLSt=<value>

This keyword specifies the display mode of the SQL command monitor. It will affect the display content of the **SQL_CMD** and **TIME_OF_SQL_CMD** columns in the **SYSUSER** system table. The SQL command monitor can contain precise or rough information for executing SQL commands. The precise information consumes more CPU time than the rough information. You can also turn off SQL command monitor to avoid CPU overhead.

default value: 1

valid range: 0 (turn off SQL command monitor)

1 (display SQL command and rough SQL command executed time)

2 (display SQL command and precise SQL command executed time)

where to use: **server side**

DB_SPDir=<string>

This keyword specifies the path for stored procedure files. The stored procedure files include the generated dynamic link library files and the entire temp files generated during stored procedure creation. A full path name must be specified for **DB_SPDIR**.

➲ **Example for DB_SPDir:**

```
DB_SPDir = /usr/DBMaker/data/spdir          ;in UNIX
```

➲ **Example for DB_SPDir:**

```
DB_SPDir = c:\dbmaker\data\spdir             ;in Windows
```

default value: (current directory which dmserver is running)

valid range: string with length < 256

see also: DB_SPInc

where to use: server side

DB_SPInc=<string>

This keyword specifies the path for stored procedure include files. It is used when extra include files are required in a generated stored procedure. A full path name must be specified for **DB_SPInc**.

➲ **Example of DB_SPInc**

```
DB_SPInc = /usr/DBMaker/data/sp\include       ;in UNIX
```

➲ **Example of DB_SPInc:**

```
DB_SPInc = c:\dbmaker\data\sp\include         ;in Windows
```

default value: (current directory which dmserver is running)

valid range: string with length < 256

see also: DB_SPDir

where to use: server side

DB_SPLog=<string>

This keyword specifies the path for stored procedure log files. The stored procedure log files include the error log files sent from the database server while creating a stored procedure and the trace log file for the stored procedure execution. A full path name must be specified for **DB_SPLog**.

➲ **Example 1, syntax for DB_SPLog:**

```
DB_SPLog = /usr/joe/mydata/splog              ;in UNIX
```

➲ **Example 2, syntax for DB_SPLog:**

```
DB_SPLog = c:\user\joe\mydata\splog          ;in Windows
```

default value: (current directory client application is running)

valid range: string with length < 256

where to use: **client side**

DB_StrOP=<value>

This keyword specifies whether space padding needs to be removed before applying the string concatenation operator (||). A value of 0 indicates space padding for fixed length CHAR type data to be kept before applying the string concatenation operator. A value of 1 indicates the space padding to be removed before applying the string concatenation operator.

The keyword can be set on the client or server. If this keyword value is not set in the **dmconfig.ini** file on the client, the option value will rely on the **dmconfig.ini** file on the server. The default value on the server is 0.

default value: 0

valid range: 0, 1

where to use: **client and server side**

DB_StrSz=<value>

This keyword indicates the length of returned data of STRING type, used only by user-defined function (UDF). Since UDFs can only return data of fixed size, these keywords can limit the size of STRING data for clients to avoid to receive too large string.

default value: 255

valid range: 1 ~ 4096

where to use: **client or server side (client has higher priority)**

DB_StSvr=<value>

This keyword is used to activate the auto update statistics server. A value of 1 indicates that the server is started. A value of 0 indicates that the server is not running. If the auto update statistics server is activated, it will recalculate database statistics daily at 3:00 AM.

default value: 1

valid range: 0 ~ 1

where to use: server side

DB_SvAdr=<string>

This keyword can be a string that specifies the TCP/IP address of the server machine or the host name of that machine. If DNS (Domain Name Service) has been set up properly on the client machine, you can even specify the domain name in this keyword. This keyword is required at connection time on all client and server machines. If this address is not correct, the connection will fail. See a network administrator or a manual on TCP/IP networking for further information.

default value: none

valid range: a.b.c.d or host (domain) name (1<=a,b,c,d <=254)

see also: DB_PtNum

where to use: **both client and server sides.**

DB_TmiFm=<string>

This keyword specifies the time input format for SQL statements. Please refer to the *ODBC Programmer's Guide*, *Appendix B* for more information.

default value: none (accept all input time formats)

valid range: **hh:mm:ss.fff**

hh:mm:ss

```
hh:mm
hh
hh:mm:ss.fff tt
hh:mm tt
hh tt
tt hh:mm:ss.fff
tt hh:mm:ss
tt hh:mm
tt hh
```

see also: **DB_TmoFm**

where to use: **client or server side (client has higher priority)**

DB_TmoFm=<string>

This keyword specifies the time output format for SQL statements. Refer to the "ODBC Programmer's Guide" for more information.

default value: **`hh:mm:ss`**

valid range: same as valid range of **DB_TmiFm**

see also: **DB_TmiFm**

where to use: **client or server side (client has higher priority)**

DB_TpFil=<string>

This keyword specifies the names of the system temp files. The file size limitation is 8TB. Users may specify up to eight system temporary files.

default value: database name with the file extension ***.TMP***.

valid range: up to eight strings with any length, separated by one comma followed by one space

see also: **DB_DbFil, DB_BbFil**

where to use: **server side**

DB_Turbo=<value>

This keyword indicates that DBMaker will run with normal catalog cache operation. If your applications rarely modify the structure of a database, you can use **DB_TURBO** mode to speed up data access. See "Performance Tuning" for more information. This keyword is used when the database is started.

default value: 0

valid range: 0, 1

where to use: **server side**

DB_UMode=<value>

This keyword indicates the user mode. A value of 1 indicates multi-user mode and a value of 0 indicates single-user mode. This keyword has no effect on DBMaker single-user programs such as ***dmsqls***, because such programs can only start a database in single-user mode. For multi-user programs, you can set this keyword to start a database in single-user or multi-user modes. This value is required when starting a database.

default value: 1

valid range: 1 (multi-user mode)

0 (single-user mode)

see also: **DB_SMode, DB_MaxCo**

where to use: **server side**

DB_UsrBb=<string>

This keyword, a special user-defined filename, specifies the physical name of the default user BLOB file used by the operating system.

➲ **Example 1, to specify the default user BLOB file name with 20 frames:**

```
[MY_DB]                                         ;database name
DB_UsrBb = /disk1/usr/f1.bb  20                 ;blob file
```

default value: database name with the extension **.BB** and a size of 2 frames.

valid range: string with length < 256

see also: **DB_BbFil**, **DB_DbDir**, **DB_DbFil**, **DB_UsrDb**, User-defined filename

where to use: **server side**

DB_UsrDb=<string>

This keyword, a special user-defined filename, specifies the physical name of the default user data file used by the operating system.

➲ **Example 1, to specify the default user data file name with 200 pages:**

```
[MY_DB]                                         ;database name
DB_USRDB = /disk1/usr/f1.db  200                ;data file
```

default value: database name with the file extension ***.DB** and* 150 pages. For example, ***db.DB***.

valid range: string with length < 256

see also: **DB_BbFil**, **DB_DbDir**, **DB_DbFil**, **DB_UsrBb**, User-defined filename

where to use: **server side**

DB_UsrFo=<string>

This keyword indicates whether user file objects can be inserted in a database. Setting this value to 1 enables the user file object function. Refer to the chapter on "Large

Object Management" for more information. This value is required when starting a database.

default value: 0

valid range: 0, 1

see also: **DB_FoDir**

where to use: **server side**

DB_UsrId=<string>

This keyword specifies the default user ID used to login at database startup or connection time.

default value: null string

valid range: string with length = or < 32

see also: **DB_PasWd**

where to use: **client side**

DD_CTimO=<value>

This keyword specifies the connect time-out value while connecting to a remote database during a DDB operation. In the DDB environment, the coordinator database server may need to establish distributed connections to remote databases.

default value: 5 (seconds)

valid range: 1 or higher

see also: **DD_DDBMd, DD_LTimO, DB_CTimO**

where to use: **server side**

DD_DDBMd=<value>

This keyword specifies whether the DDB (Distributed DataBase) function is enabled on the database server. Turn the value on in order to use DDB operations or table replication functions.

default value: 0

valid range: 0, 1

see also: **DD_GTSvr**

Where to use: server side

DD_GTItv=<string>

This keyword specifies the schedule of the GTRECO demon to solve pending global transactions. It is used only when the GTRECO server is on.

The input format is 'D hh:mm:ss'.

default value: 00:10:00

valid range: 0 days ~ 24855 days

see also: **DD_GTSvr**

where to use: server side

DD_GTSVR=<value>

This keyword specifies whether to start up the GTRECO (Global Transaction RECOvery) demon while DDB mode is on. The GTRECO demon will automatically solve the pending global transactions that cross the DBMaker database servers.

default value: 1

valid range: 0, 1

see also: **DD_DDBmd**, **DD_GTItv**

where to use: **server side**

DD_LTimO=<value>

This keyword specifies the lock timeout value for the distributed data access during the DDB operation. It takes effect on server-to-server data access only. Refer to **DB_LTimO** for more information on the timeout value .

default value: 5 (seconds)

valid range: -1 or higher

see also: **DD_DDBmd**, **DD_CTimO**, **DB_LTimO**

where to use: server side

DM_DifEn=<value>

This keyword needs to be set under the global section in **DM_COMMON_OPTION,** to denote whether or not to allocate a new environment handle when requested. The global section in **DM_COMMON_OPTION,** in the **dmconfig.ini** file, is for global settings across databases. Keywords like the DM_DIFEN, apply to all databases in the **dmconfig.ini** file when it is specified. Do not change it unless advice from DBMaker customer support for a special case arises.

default value: 1

valid range: 0, 1

where to use: client side

LG_NPFun=<string>

This keyword is set only in the **DM_COMMON_OPTION** section and specifies the unlogged functions. Its value is an empty string or some ODBC function names separated by commas (,). This keyword is valid only if **LG_PTFun** is not defined in the **dmconfig.ini** file. Once this keyword is specified, the functions listed in the string will not be logged.

default value: “” (empty string, all functions will be logged)

valid range: **function list string** (eg. “SQLError, SQLGetDiagRec”)

see also: **LG_Path**, **LG_PTFun**, **LG_Trace**, **LG_Time**

where to use: **client side**

LG_Path=<string>

This keywordis only set in the **DM_COMMON_OPTION** section and specifies the file path name of the output log file.

default value: c:\odbclog.log (for Win32), ./odbclog.log (for UNIX)

valid range: string with length < 256

see also: **LG_NPFun**, **LG_PTFun**, **LG_Time**, **LG_Trace**

where to use: **client side**

LG_PTFun=<string>

This keyword is set in the **DM_COMMON_OPTION** section and specifies the logged functions. Its value is an empty string or some ODBC function names, separated by commas (,). Once this keyword value is specified, only the functions listed in the string will be logged. If **LG_PTFun** and **LG_NPFun** are set, only the **LG_PTFun** will take effect.

default value: none (all functions will be logged)

valid range: **function list string** (eg. “SQLError, SQLGetDiagRec”)

see also: **LG_NPFun, LG_Path, LG_Time, LG_Trace**

where to use: **client side**

LG_Time=<value>

This keyword is set in the **DM_COMMON_OPTION** section and specifies whether to log time spent on each function. Setting this value to 1 logs time spent for each function in the output log file, setting it to 0 does not log time. This information can help users find performance bottlenecks in an ODBC program.

default value: 0

valid range: 0, 1

see also: **LG_NPFun**, **LG_Path**, **LG_PTFun**, **LG_Trace**

where to use: **client side**

LG_Trace=<value>

This keyword is set in the **DM_COMMON_OPTION** section and specifies whether the ODBC log is turned on or off. Setting this value to 1 will turn on the ODBC log and setting it to 0 will turn it off. When the ODBC log is on, the called ODBC function, input parameters, output parameters, and returned code or error information will be logged to a specified file. See the following keywords for more detailed information.

default value: 0

valid range: 0, 1

see also: **LG_NPFun, LG_Path, LG_PTFun, LG_Time**

where to use: **client side**

RP_BTime=<value>

This keyword is used for database replication and specifies the starting time of database replication. The format is YYYY/MM/DD hh:mm:ss, e.g. 2000/1/1 01:30:00. You can use **RP_Iterv** to specify the schedule of the database replication.

default value: time of the primary database starting

Valid range: **YYYY/MM/DD hh:mm:ss** (e.g. 2000/1/1 00:00:00)

See also: **DB_SMode**, **RP_Clear**, **RP_Iterv**, **RP_Primy**, **RP_PtNum**, **RP_ReTry**, **RP_SlAdr.**

Where to use: **server side of the primary database.**

RP_Clear=<value>

This keyword is used for database replication and specifies whether DBMaker should clear the backup files after replicating them to remote databases. The value 1 clears the files, 0 will keep them.

Default value: **0** (no).

Valid range: **0** (no), **1** (yes).

See also: **DB_SMode**, **RP_BTime**, **RP_Iterv**, **RP_Primy**, **RP_PtNum**, **RP_ReTry**, **RP_SlAdr**.

Where to use: **server side of the primary database.**

RP_LgDir=<string>

This keyword, used for the asynchronous table replication, specifies the directory where DBMaker puts replication log files for the asynchronous table replication. The replication log files are binary and users should not manually remove them.

Default value: the subdirectory named **TRPLOG** under the database home directory

Valid range: string with length < 256

See also: **DB_AtrMd**, **DB_EtrPt**

Where to use: **server side of the primary database**

RP_Iterv=<value>

This keyword specifies the schedule for the database replication. The format is dd-hh:mm:ss, e.g. 1-12:00:00; one day and 12 hours. You can use **RP_BTime** to specify the starting time of database replication.

Default value: **1-00:00:00** (one day).

Valid range: **0** days ~ **24855** days.

See also: **DB_SMode**, **RP_BTime**, **RP_Clear**, **RP_Primy**, **RP_PtNum**, **RP_ReTry**, **RP_SlAdr.**

where to use: **server side of the primary database**

RP_Primy=<string>

This keyword, used for the database replication, specifies the address of the primary database.

Default value: none

Valid range: a,b,c,d or host (domain) name (1<=a,b,c,d <=254)

See also: **DB_SMode**, **RP_BTime**, **RP_Clear**, **RP_Iterv**, **RP_PtNum**, **RP_ReTry**, **RP_SlAdr.**

Where to use: **server side of the slave database**

RP_PtNum=<value>

This keyword, used for the database replication, specifies the port number of the RP_RECV_SERVER daemon the slave database. It must be different from the **DB_PtNum** used by the slave database and the same as the port number specified in **RP_SlAdr** on the primary database.

Default value: **23001**

Valid range: 1024~65535

See also: **DB_SMode**, **RP_BTime**, **RP_Clear**, **RP_Iterv**, **RP_Primy**, **RP_ReTry**, **RP_SlAdr**

Where to use: **server side of the slave database.**

RP_Reset=<value>

RP_RESET is used to indicate to DBMaker to reset the asynchronous table replication system while starting a database. If RP_RESET is set to 1, DBMaker will clear all unsent asynchronous table replication logs, remove all **.TRP** files in the **RP_LgDir** (default is DB_DBDIR/TRPLOG) directory, and reset the asynchronous table replication status while starting the database. That is, DBMaker will ignore all the unsent asynchronous table replication data. After starting database, DBMaker server will reset the RP_RESET value to 0 to prevent duplicate resetting asynchronous table replication in next time of starting database.

Default value: 0

Valid range: 0-1

See also: **RP_LgDir**.

Where to use: **server side of primary database.**

RP_ReTry=<value>

This keyword, used for the database replication, specifies how many times DBMaker will try to connect to remote databases after a network failure.

default value: **0**

valid range: **0 ~**

see also: **DB_SMode**, **RP_BTime**, **RP_Clear**, **RP_Iterv**, **RP_Primy**, **RP_PtNum**, **RP_SlAdr**

where to use: **server side of the primary database**

RP_SlAdr=<string>

This keyword is used for database replication and specifies the slave databases for the primary database to send data. DBMaker supports 1 to 8 slave databases for each primary.

➲ **Syntax for RP_SlAdr:**

```
RP_SlAdr = address[:port number] {, address[:port number]}
```

The default port number is 23001. In slave databases, commas or blank spaces can separate the information.

➲ **Example, RP_SlAdr port numbers:**

```
RP_SlAdr = 192.168.9.222:5100, Server2:5101, Server3
```

There are three slave databases. One is 192.168.9.222 with port number 5100, another is Server2 with port number 5101, and the other is Server3 with default port number 23001.

default value: none.

See also: **DB_SMode**, **RP_BTime**, **RP_Clear**, **RP_Iterv**, **RP_Primy**, **RP_PtNum**, **RP_ReTry**

Where to use: **server side of the primary database.**

User-defined filename=<physical filename> <pages>

DBMaker allows users to create data or BLOB files and add them to a tablespace when the original tablespace has been filled. Users generally specify a logical file name without a full path when creating a file. Users can map the logical file name to a physical file path name that is used by the operating system to access the file.

➲ **Mapping a file in the dmconfig.ini file:**

```
FILE1 = /disk1/usr/datafile 100
```

Although you specify **FILE1** in DBMaker, DBMaker will create a file **/disk1/usr/datafile** with 100 pages (4KB per page). If this file has to be moved to another directory, change the physical file name in the **dmconfig.ini** file. There is no need for modification to your programs or SQL scripts. The rule for **DB_DbDir** also applies to user-defined files.

valid range: file name — string with length < 256

size — 2 ~ 524287

see also: **DB_DbDir**, **DB_UsrBb**, **DB_UsrDb**

where to use: **server side**

B System Catalog Reference

Part of the definition for a relational database is that all database information must be represented at the logical level in the same way as user data. This information is stored in the *system catalog*, allowing authorized users to use SQL to access information on the database in the same way they access data in SQL tables. This appendix contains descriptions of the system catalog tables and views, organized alphabetically by name. You can query these system catalog tables to view detailed status of a database.

B.1 The System Catalog

The system catalog is a set of tables that contains information on all objects in the database. The system catalog is also known as the *data dictionary*.

All system catalogs are owned by SYSTEM, and can be read by any user that has at least the Connect Authority. Since system catalogs belong to SYSTEM, you cannot DROP a system table or a system-defined column, and you cannot INSERT or DELETE rows in a system table.

B.2 DBMaker System Catalog Tables

The following table lists all of the system catalog tables in DBMaker, and a brief description of what is contained in each table.

TABLE NAME	CONTENTS
SYSAUTHCOL	Column privilege information
SYSAUTHEXE	Executable object privilege information
SYSAUTHGROUP	Group information
SYSAUTHMEMBER	Group member information
SYSAUTHTABLE	Table privilege information
SYSAUTHUSER	Security level information
SYSCMDINFO	Stored command information
SYSCOLUMN	Column information
SYSCONINFO	Connection information
SYSDBLINK	Database link information
SYSDOMAIN	Domain information
SYSFILE	File information
SYSFILEOBJ	File object information
SYSFOREIGNKEY	Foreign key information
SYSGLBTRANX	DDB global transaction information
SYSINDEX	Index information
SYSINFO	Database system information
SYSIOINFO	I/O Information
SYSLOCK	Lock information.
SYSOPENLINK	Open link information
SYSPENDTRANX	Pending distributed transaction information

SYSPROCINFO	Stored procedure information
SYSPROCPARAM	Stored procedure parameter information
SYSPROJECT	ESQL project information
SYSPUBLISH	Table replication source information
SYSSUBSCRIBE	Table replication destination information
SYSSYNONYM	Synonym information
SYSTABLE	Table information
SYSTABLESPACE	Tablespace information
SYSTEXTINDEX	Text index information
SYSTRIGGER	Trigger information
SYSTRPDEST	Table replication information
SYSTRPJOB	Information for recording all jobs to be replicated
SYSTRPPOS	Information for Distributor to prune transaction log files
SYSUSER	Information on users logged into the database
SYSUSERFUNC	User-defined function information
SYSVIEWDATA	View information
SYSWAIT	Waiting connection information

SYSAUTHCOL

The SYSAUTHCOL table lists the columns in all tables on which a user has been granted object privileges. If a user is allowed to perform some operations such as INSERT, UPDATE or REFERENCE on the authorized table with all columns (i.e. the return value of INS_ALL, UPD_ALL, or REF_ALL in SYSAUTHTABLE is 1), then the return values of the columns, INS, UPD, REF in SYSAUTHCOL should be ignored.

Column Name	Description
COLUMN_NAME	Name of the column on which privileges have been granted.
TABLE_NAME	Name of the table the column belongs to.
GRANTEE	Name of the user granted privileges on the column. Must be a valid user or group name.
TABLE_OWNER	Name of the user who created the table.
INS	1 — User has the privilege to insert data into the specified column. 0 — User does not have the privilege to insert data into the specified column.
UPD	1 — User has the privilege to update data in the specified column. 0 — User does not have the privilege to update data in the specified column.
REF	1 — User has the privilege to create a constraint that refers to the specified column. 0 — User does not have the privilege to create a constraint that refers to the specified column.

SYSAUTHEXE

The SYSAUTHEXE table contains executable object information.

Column Name	Description
OBJNAME	Name of the executable object.
OWNER	User who created the executable object.
OBJTYPE	Type of executable object, such as "Procedure", "Command", "Project", etc.
GRANTEE	Name of the user granted privileges on the executable object.

SYSAUTHGROUP

The SYSAUTHGROUP table gives the names of all valid groups in the database.

COLUMN NAME	DESCRIPTION
GROUP_NAME	Name of the group
GROUP_OWNER	User who created the group
NUM_MEMBERS	Number of members in this group

SYSAUTHMEMBER

The SYSAUTHMEMBER table lists all members who belong to a group.

COLUMN NAME	DESCRIPTION
MEMBER_NAME	Name of the member who belongs to the group
GROUP_NAME	Name of the group

SYSAUTHTABLE

The SYSAUTHTABLE table is a list of all privileges which have been granted on tables, and who they were granted to.

COLUMN NAME	DESCRIPTION
TABLE_NAME	Name of the table or view on which privileges have been granted.
GRANTEE	Name of the user granted privileges on the table.
TABLE_OWNER	User who created the table or view.
NUM_RPI_COLS	Number of columns that have privileges granted on them in the table or view.

Column Name	Description
SEL_ALL	1 — User has the privilege to select data from all columns in the specified table or view. 0 — User does not have the privilege to select data from columns in the specified table or view.
DEL_ALL	1 — User has the privilege to delete data in all columns in the specified table or view; 0 — User does not have the privilege to delete data from columns in the specified table or view.
INS	1 — User has the privilege to insert data into specific columns in the specified table or view. 0 — User does not have the privilege to insert data into any columns in the specified table or view.
INS_ALL	1 — User has the privilege to insert data into all columns in the specified table or view. 0 — User does not have the privilege to insert data into all columns in the specified table or view, but may still have privileges on individual columns (see INS).
UPD	1 — User has the privilege to update data in specific columns in the specified table or view. 0 — User does not have the privilege to update data in any columns in the specified table or view.
UPD_ALL	1 — User has the privilege to update data in all columns in the specified table or view. 0 — User does not have the privilege to update data in all columns in the specified table or view, but may still have privileges on individual columns (see UPD).
ALT_ALL	1 — User has the privilege to alter the definition of the specified table or view. 0 — User does not have permission to alter the definition of the specified table or view.

COLUMN NAME	DESCRIPTION
IDX_ALL	1 — User has the privilege to create or drop indexes on the specified table or view. 0 — User does not have the privilege to create or drop indexes on the specified table or view.
REF	1 — User has the privilege to create a CONSTRAINT, WHICH refers to specific columns in the specified table or view. 0 — User does not have the privilege to create a constraint on any columns in the specified table or view.
REF_ALL	1 — User has the privilege to create a CONSTRAINT, WHICH refers to columns in the specified table or view. 0 — User does not have the privilege to create a constraint which refers to columns in the specified table or view, but may still have privileges on individual columns (see REF).

SYSAUTHUSER

The SYSAUTHUSER table lists the names and authority levels of valid users in a database.

COLUMN NAME	DESCRIPTION
USER_NAME	User ID of each valid user in the database; a user is considered valid if they have been granted CONNECT authority.
DBA	**1** — User has DBA Authority. **0** — User does not have Dba authority.
RESOURCE	**1** — User has resource authority. **0** — User does not have resource authority.

SYSCMDINFO

The SYSCMDINFO table contains stored command information.

Column Name	Description
MODULENAME	Module name
CMDNAME	Command name
CMDOWNER	Command owner
STATEMENT	Command statement
NUM_PARM	Number of parameters
STATUS	1: valid; 0: invalid

SYSCOLUMN

This table lists every column for each table and view, including the columns in the system catalog tables. In the SCALE and RADIX columns, a value of -1 is returned where SCALE and RADIX are not applicable to the data type found in that column.

Column Name	Description
COLUMN_NAME	Name of the column
TABLE_NAME	Name of the table that owns this column
TABLE_OWNER	Name of the user who created the table
COLUMN_ORDER	Order of the column in the table
NULLABLE	**1** — Column allows null values **0** — Column does not allow null values
TYPE_NAME	Type name of the column. Can be any of the following: BINARY, CHAR, NCHAR, DATE, DECIMAL, DOUBLE, FILE, FLOAT, INTEGER, LONG VARCHAR, NCLOB, LONG VARBINARY, SERIAL, SMALLINT, TIME, TIMESTAMP, VARCHAR, or NVARCHAR

COLUMN NAME	DESCRIPTION
PRECISION	Precision of the column
SCALE	Scale of the column
RADIX	Radix of the column
ASCII_DEF	Default value of ASCII form for the column
CONSTRAINT	Constraint for the column
REMARKS	A description of the column

SYSCONINFO

The SYSCONINFO table contains information about the connection of the database.

COLUMN NAME	DESCRIPTION
CONNECTION_ID	Connection ID
LAST_SERIAL	The last serial number working on the updated column, that is of a SERIAL type
LAST_OID	The object ID (OID) of the last inserted record.
INFO1	Reserved
INFO2	Reserved
INFO3	Reserved
INFO4	Reserved

SYSDBLINK

The SYSDBLINK table contains information on remote database links.

COLUMN NAME	DESCRIPTION
OWNER	Link Owner
DB_LINK	Link Name

Column Name	Description
DBSVR	A database section, which contains the remote database information.
USER_NAME	User Name in the remote database

SYSDOMAIN

The SYSDOMAIN table contains information on domains created in the database.

Column Name	Description
DOMAIN_NAME	The name of the domain
DOMAIN_OWNER	The name of the user who created the domain
ASCII_DEF	Default value of ASCII form for the domain
TYPE_NAME	Type name of the column. Can be any of the following: BINARY, CHAR, NCHAR, DATE, DECIMAL, DOUBLE, FILE, FLOAT, INTEGER, LONG VARCHAR, NCLOB, LONG VARBINARY, SERIAL, SMALLINT, TIME, TIMESTAMP, VARCHAR, or NVARCHAR
DATA_LEN	The size of the data type for the domain
PRECISION	Precision of the domain
SCALE	Scale of the domain
CONSTRAINT	Constraint of the domain

SYSFILE

The SYSFILE table contains information on files in the database.

Column Name	Description
FILE_NAME	Logical file name

COLUMN NAME	DESCRIPTION
FILE_TYPE	File type: 1 (data file) or 2 (BLOB file)
TS_NAME	Name of the tablespace the file is in
FILE_NPAGES	Number of pages in the file. If the tablespace is an AUTOEXTEND tablespace, then FILE_NPAGES may be less than the physical FILE_NPAGES in the file.
RAWDEV_OFFSET	Not supported in this version of DBMaker

SYSFILEOBJ

The SYSFILEOBJ table contains information on the file objects in the database. This includes both system and user file objects.

COLUMN NAME	DESCRIPTION
FILE_TYPE	**00** — system file object **01** — user file object
SHARE	The number of records that share the file object
FILE_NAME	The full path with the filename to show where the file object is located.

SYSFOREIGNKEY

The SYSFOREIGNKEY table contains information on all foreign keys in the database.

COLUMN NAME	DESCRIPTION
FK_TBL_NAME	Name of the child table (the table the foreign key is defined on)
PK_TBL_NAME	Name of the parent table for the foreign key
FK_TBL_OWNER	Owner of the child table
PK_TBL_OWNER	Owner of the parent table

Column Name	Description
FK_NAME	Name of the foreign key
UPD_ACT	Update referential action 0 — No action 1 — Set NULL 2 — Cascade 3 — Set default value
DEL_ACT	Delete referential action 0 — No action 1 — Set null 2 — Cascade 3 — Set default value

SYSGLBTRANX

The SYSGLBTRANX table contains information on global transactions.

COLUMN NAME	DESCRIPTION
STATE	Global transaction state: 0 (ISSUE) — A transaction branch is issued, but the participant has not prepared all participants. 1 (PREPARE) — The participant is prepared, but waits for it's father participant to decide whether to commit or abort. 2 (COMMIT) — The participant has decided to commit the global transaction. 3 (PEND_TO_COMMIT) — After crash recovery, this transaction branch has been added to the commit queue and is waiting to be committed. 4 (PEND_TO_ABORT) — After crash recovery, this transaction branch has been added to the abort queue and is pending an abort.
PARTICIPANT	Global transaction participant
GLBTRANXID	Global transaction ID

SYSINDEX

The SYSINDEX table contains information on indexes in the database. A -1 in the NUM_PAGE, NUM_LEVEL, NUM_LEAF, DIST_KEY, NUM_PAGE_KEY, or CLSTR_COUNT columns means those values are not applicable.

COLUMN NAME	DESCRIPTION
INDEX_NAME	Name of the index
TABLE_NAME	Name of the table that the index is defined on
TABLE_OWNER	Owner of the table that the index is defined on
TS_NAME	Specifies the tablespace that the index is stored on

COLUMN NAME	DESCRIPTION
UNIQUE	Status flag to indicate uniqueness of the index: **0** — non-unique **1** — unique **3** — primary key
NUM_COL	Number of columns in the index.
NUM_PAGE	Number of index pages.
NUM_LEVEL	Number of levels.
NUM_LEAF	Number of leaf pages.
DIST_KEY	Number of distinct keys.
NUM_PAGE_KEY	Number of pages per key.
CLSTR_COUNT	Cluster count; the number of page I/O while we use the index to access data page. It is related to the number of buffers.

SYSINFO

The SYSINFO table gives information on the current state of a database.

The SYSINFO table's schema is different from other system tables. The schema is given below:

```
SYSTEM.SYSINFO (char(4) ID, varchar(32) INFO, varchar(32) VALUE);
```

The meaning of each column follows:

- **ID**: Item identifier. The system information is cataloged according to this ID. The first two characters represent the category, and the following 2 characters identify the item within the category. E.g. For the ID '0105' representing NUM_LOGICAL_READ, the '01' means it belongs to the page and I/O category, and the '05' is its sequence in the page and I/O category. Users can sort or filter the SYSINFO by ID.
- **INFO**: Item name of the system information. E.g. the name 'NUM_LOGICAL_READ' means the number of logical disk read.

- **VALUE**: All values of system information are returned as VARCHAR data.

➲ **Example**

The statement below will show the number of logical disk read:

```
dmSQL> select INFO, VALUE from SYSTEM.SYSINFO where INFO = 'NUM_LOGICAL_READ';
ID      INFO                               VALUE
====    ==================
        =========================================
0105    NUM_LOGICAL_READ          338

1 rows selected
```

Below lists all items in the SYSINFO catalog.

PAGE AND I/O INFORMATION:

SYSINFO.ID	SYSINFO.INFO	DESCRIPTION
0101	NUM_IDX_SPLIT	Number of index page splits occurring.
0102	NUM_PAGE_COMPRESS	Number of data pages compressed, i.e. page reorganization.
0103	NUM_PHYSICAL_READ	Number of physical disk reads. I/O unit is Page.
0104	NUM_PHYSICAL_WRITE	Number of physical disk writes. I/O unit is Page.
0105	NUM_LOGICAL_READ	Number of logical reads. I/O unit is Page.
0106	NUM_LOGICAL_WRITE	Number of logical writes. I/O unit is Page.
0107	NUM_PAGE_BUF	Number of page buffers. Counting unit is Page.

NOTE *1 Page = 4096 bytes.*

JOURNAL INFORMATION:

SYSINFO.ID	SYSINFO.INFO	DESCRIPTION
0201	NUM_JNL_BLK_READ	Number of journal blocks read from journal files expressed in blocks.
0202	NUM_JNL_BLK_WRITE	Number of journal blocks written to journal files expressed in blocks.
0203	NUM_JNL_REC_WRITE	Number of journal records generated. The new journal record is placed in the journal buffer first.
0204	NUM_JNL_FRC_WRITE	Number of journal forced writes. This number is the I/O number of flushing dirty journal buffer to disk.
0205	NUM_JOURNAL_FILE	Number of journal files.
0206	NUM_JOURNAL_BLOCKS	Number of journal blocks in a file. The total number of journal blocks in a database is: NUM_JOURNAL_FILE* NUM_JOURNAL_BLOC KS
0207	NUM_JNR_BLOCK_FREE	Number of free journal blocks.
0208	CURRENT_JOURNAL_FN	The file number of the currently used journal file
0209	CURRENT_JOURNAL_BN	The current block number of the journal file. Each journal block of journal files has a unique address that is formed by CURRENT_JOURNAL_F

SYSINFO.ID	SYSINFO.INFO	DESCRIPTION
		N and CURRENT_JOURNAL_BN. The block number of a journal file is counted from 0.

NOTE *1 Block = 512 bytes.*

TRANSACTION INFORMATION:

SYSINFO.ID	SYSINFO.INFO	DESCRIPTION
0301	NUM_STARTED_TRANX	Number of started transactions.
0302	NUM_COMMITED_TRANX	Number of committed transactions.
0303	NUM_ABORTED_TRANX	Number of aborted transactions.
0304	NUM_CHECKPOINT	Number of checkpoints.
0305	NUM_COMMIT_WAITER	Number of transactions awaiting group commit.

LOCK INFORMATION:

SYSINFO.ID	SYSINFO.INFO	DESCRIPTION
0401	NUM_ROW_LOCK_UPG	Number of page locks escalated (i.e. row locks escalated to a page lock).
0402	NUM_PAGE_LOCK_UPG	Number of table locks escalated (i.e. page locks escalated to a table lock).
0403	NUM_LOCK_TIMEOUT	Number of failed locks due to timeout.

SYSINFO.ID	SYSINFO.INFO	DESCRIPTION
0404	NUM_LOCK_WAIT	Number of locks waiting.
0405	NUM_LOCK_REQUEST	Number of locks requested.
0406	NUM_DEADLOCK	Number of deadlocks detected.

CONNECTION INFORMATION:

SYSINFO.ID	SYSINFO.INFO	DESCRIPTION
0501	NUM_MAX_HARD_CONNECT	Maximum number of allowed connections for a database (hard limitation of connection, i.e. **DB_MaxCo** when database start with new journal or creating a new database).
0502	NUM_MAX_SOFT_CONNECT	Maximum number of allowed connections at a time (soft limitation of connection, i.e. **DB_MaxCo** when database started normally). The soft limitation of connections is less or equal to the hard limitation of connections. (*In previous versions of DBMaker this was called NUM_MAX_TRANX*)
0503	NUM_CONNECT	Number of currently active connections. (*In previous versions of DBMaker this was called NUM_ACT_TRANX*)
0504	NUM_PEAK_CONNECT	Maximum number of active connections at a time (peak number of active connections).

Data Operation Information:

SYSINFO.ID	SYSINFO.INFO	Description
0601	NUM_SQL_SELECT	Number of SELECT operations.
0602	NUM_SQL_INSERT	Number of INSERT (including INSERT INTO) operations.
0603	NUM_SQL_UPDATE	Number of UPDATE operations.
0604	NUM_SQL_DELETE	Number of DELETE operations.
0605	NUM_SQL_PREPARE	Number of SQLPrepare() calls to server.
0606	NUM_SQL_EXECUTE	Number of SQLExecute() calls to server.
0607	NUM_SQL_EXEDIRECT	Number of SQLExecDirect() calls to server.
0608	NUM_SQL_FETCH	Number of fetched data passing across the network.

Database Information:

SYSINFO.ID	SYSINFO.INFO	Description
0701	SYSINFO_RESET_TIME	Time the counter of SYSINFO was restarted (*new*) This is used for record the time the SYSINFO is reset. The setting happen when: 1. dmSQL> set SYSINFO clear; 2. One counter is overflow, and then reset all counters. The checking is done when:

SYSINFO.ID	SYSINFO.INFO	DESCRIPTION
		2-1. Each time to select SYSINFO table. 2-2. Every about 5 seconds by I/O Server.
0702	DCCA_SIZE	Total size of DCCA. **Byte unit.**
0703	FREE_DCCA_SIZE	Available size of DCCA. **Byte unit.**
0704	DDB_MODE	Distributed database mode; **ON**: Enable; **OFF**: Disable.
0705	BACKUP_MODE	Backup mode; . **NON-BACKUP**: non backup mode (DB_BMode = 0) . **BACKUP-DATA**: backup data only mode (DB_BMode = 1) . **BACKUP-DATA-AND-BLOB**: backup data and BLOB mode (DB_BMode = 2)
0706	USER_FO_MODE	User file object mode; **ON**: Enable; **OFF**: Disable.
0707	READ_ONLY_MODE	Read-only mode; **ON**: Enable; **OFF**: Disable.
0708	FRAME_SIZE	BLOB frame size. **Byte unit.**
0709	CREATE_DB_TIME	Time the database was created.
0710	START_DB_TIME	Time the database was started.
0711	VERSION	DBMaker version.
0712	FILE_VERSION	Database file version.

SYSINFO.ID	SYSINFO.INFO	DESCRIPTION
0713	FORCE_NEW_JNL_TIME	Time the database startup with force new journal.
0714	START_NO_JNL_TIME	Time of turning the journal off.
0715	END_NO_JNL_TIME	Time of turning the journal on.

SYSTEM INFORMATION:

SYSINFO.ID	SYSINFO.INFO	DESCRIPTION
0801	CPU_USAGE	The average CPU load during a short period (about 5 seconds) (*new*). (**0 ~ 100** %) SUPPORTED PLATFORM: Solaris, LINUX, Windows 2000 (only count the first CPU, and need pdh.dll library). You must start I/O Server to enable this item.
0802	TOTAL_MEMORY	Total physical memory. **Byte unit**. SUPPORTED PLATFORM: Solaris, LINUX, Windows NT/2000, UNIX that supports POSIX standard.
0803	TOTAL_FREE_MEMORY	The current free physical memory (*new*). **Byte unit**. SUPPORTED PLATFORM: Solaris, LINUX, Windows NT/2000, UNIX that supports POSIX standard.
0804	TOTAL_SWAP_SPACE	Total swap space. **Byte unit.** SUPPORTED PLATFORM: Solaris, LINUX, Windows

SYSINFO.ID	SYSINFO.INFO	DESCRIPTION
		NT/2000.
0805	TOTAL_FREE_SWAP_ SPACE	The current free swap space. (*new*). **Byte unit.** SUPPORTED PLATFORM: Solaris, LINUX, Windows NT/2000.

For unsupported versions the value is NULL. The SYSINFO catalog is the collection of accumulated counters.. Users should know that there are two ways to reset SYSINFO:

1. A user executes the "set SYSINFO clear" command.
2. When one counter overflows, all counters in the SYSINFO will be reset. DBMaker checks for overflow:
   - **a)** Each time the SYSINFO table is selected.
   - **b)** Every 5 seconds by the I/O Server.

Users can get the reset time by executing the following statement:

```
dmSQL> select VALUE from SYSTEM.SYSINFO where INFO = 'SYSINFO_RESET_TIME';
```

SYSLOCK

The SYSLOCK table contains information on status for locks on objects in s database.

NOTE *Lock granularity can be SYSTEM, TABLE, PAGE, or TUPLE, lock status can be GRANTED, WAITING, or CONVERT, and lock mode can be NONE, IS, S, IX, SIX, or X.*

COLUMN NAME	DESCRIPTION
LK_OBJECT_ID	OID of locked object.
TABLE_ID	OID of the table containing the locked object.
LK_GRAN	Lock granularity, SYSTEM, TABLE, PAGE, or TUPLE.

HOLD_LK_CONNECTION	Connection ID that is holding the lock on the object.
LK_STATUS	Lock status, GRANTED, WAITING, or CONVERT.
LK_CURRENT_MODE	Current lock mode of object.
LK_NEW_MODE	New lock mode of object.

SYSOPENLINK

The SYSOPENLINK table contains information on open database links.

COLUMN NAME	DESCRIPTION
DB_LINK	Open link.
DBSVR	Server.
USER_NAME	User name.
TXN_STATUS	Transaction status. 'R' — Read 'W' — Write 'N' — No transaction

SYSPENDTRANX

The SYSPENDTRANX table contains information on uncommitted transactions in a distributed database environment.

COLUMN NAME	DESCRIPTION
XIDFORMAT	Format ID denoting the type of the global coordinator, DBMaker is 22873.
PREPAREDTIME	Time of the prepared commit transaction.
GLBTRANXID	Global transaction ID.

SYSPROCINFO

The SYSPROCINFO table contains information on stored procedures.

Column Name	Description
QUALIFIER	Qualifier.
PROC_OWNER	Procedure owner.
NAME	Procedure name.
NUM_INPUT_PARAMS	Number of input parameters.
NUM_OUTPUT_PARAMS	Number of output parameters.
NUM_RESULT_SETS	Number of result sets.
REMARKS	Remarks.
PROC_TYPE	Procedure type. 1 (SQL_PT_PROCEDURE) — Procedure 2 (SQL_PT_FUNCTION) — Function

SYSPROCPARAM

The SYSPROCPARAM table contains information on stored procedure parameters.

Column Name	Description
QUALIFIER	Qualifier.
OWNER	Procedure owner.
PROC_NAME	Procedure name.
PARAM_NAME	Parameter name.
PARAM_TYPE	Parameter type. 1 (SQL_PARAM_INPUT) — Input. 3 (SQL_PARAM_OUTPUT) — Output. 4 (SQL_RETURN_VALUE) — Return value. 5 (SQL_RESULT_COL) — Result set.

Column Name	Description
DATA_TYPE	Data type.
TYPE_NAME	Type name.
PRECISION	Precision.
LENGTH	Length.
SCALE	Scale.
RADIX	Radix.
NULLABLE	Nullable column. **1** — Allows null values. **0** — Does not allow null values.
REMARKS	Remarks.

SYSPROJECT

The SYSPROJECT table contains information on ESQL projects.

Column Name	Description
PROJECT_NAME	Project name.
PROJECT_OWNER	Project owner.
MODULE_NAME	Module name.
MODULE_OWNER	Module owner.
MODULE_SOURCE	Module source.
REF_CMD	Internal usage.

SYSPUBLISH

The SYSPUBLISH table contains information on table replication sources.

Column Name	Description
REPLICATION_NAME	Name of replication.
TYPE	S — Synchronous A — Asynchronous.
TABLE_OWNER	Owner of table being replicated.
TABLE_NAME	Name of table being replicated.
NUM_PROJECT	Number of projected columns.
FRAGMENT	Fragment string.
NUM_SUBSCRIBER	Number of subscribers.

SYSSUBSCRIBE

The SYSSUBSCRIBE table contains information on table replication targets.

Column Name	Description
BASE_TABLE_OWNER	Base table owner.
BASE_TABLE_NAME	Base table name.
REPLICATION_NAME	Replication name.
DB_LINK	Database link.
TABLE_OWNER	Table owner.
TABLE_NAME	Table name.

SYSSYNONYM

The SYSSYNONYM table contains information on synonyms defined in a database.

Column Name	Description
SNAME	Synonym name.
OWNER	Synonym owner.

TV_NAME	The source table/view name of the synonym.
TV_OWNER	Table/view owner.
TV_LINK	Link name of a table or view in a remote database.
TV_SERVER	Database name of a table or view in a remote database.

SYSTABLE

The SYSTABLE table contains information on tables in the database.

COLUMN NAME	DESCRIPTION
TABLE_NAME	Name of the table.
TABLE_OWNER	Owner of the table.
TABLE_TYPE	Table type: SYSTEM TABLE, SYSTEM VIEW, TABLE, or VIEW.
LOCKMODE	Lock mode applied to the table: **T** — table lock **P** — page lock **R** — row lock The default lock mode is page lock.
CACHEMODE	Cache mode of the full table scan: T — there is caching (true). F — there is no caching (false).
TS_NAME	Name of the tablespace the table is in.
TABLE_OID	The OID of the table.
TABLE_VERSION	
NUM_COL	Number of columns in the table.
NUM_INDEX	Number of indexes on the table.

COLUMN NAME	DESCRIPTION
NUM_PAGE	Number of pages in the table. The default value is -1. When the user updates statistics on the table, the true value of NUM_PAGE will be returned.
NUM_FRAME	Number of BLOB frames in the table
NUM_ROW	Number of rows in the table. The default value is -1. When the user updates statistics on the table, the true value of NUM_ROW will be returned.
NUM_INDIRECT_ROW	Number of indirect rows
AVERAGE_LENGTH	Average length for data in the table. Default value is -1. When the user updates statistics on the table, the true value of AVERAGE_LENGTH will be returned.
CREATE_TIME	Time that the table was created.
UPD_STS_TIME	Last time table's statistics were updated.
UPD_STS_INTERVAL	Update statistics time interval.
CONSTRAINT	Constraint for the table.
FILLFACTOR	The **FILLFACTOR** specifies the percentage of the page that can be filled before it stops allowing new data to be inserted (to allow room for updates). The default value is 100 (%).
SERIAL_COL_ID	Serial column is located in the n^{th} column in the table.
SERIAL_START_NO	Serial column starting number. The default value is 1.
REMARKS	A description of the table.
NUM_TRIG	Number of triggers on table.
NUM_TEXTINDEX	Number of text indexes on table.
NUM_PUBLICATION	Number of publications on table.

COLUMN NAME	DESCRIPTION
NUM_DEST	Number of target databases for asynchronous replication.

SYSTABLESPACE

The SYSTABLESPACE table contains information on all tablespaces in the database.

COLUMN NAME	DESCRIPTION
TS_NAME	Name of the tablespace.
TS_TYPE	Type of tablespace: **1** — AUTOEXTEND. **0** — NORMAL.
NUM_FILES	Number of files in the tablespace.
NUM_PAGES	Number of pages in the tablespace. If the tablespace is autoextend, then NUM_PAGES may be less than the real NUM_PAGES in the tablespace.
NUM_FREE_PAGES	Number of free pages available in the tablespace.
NUM_FRAMES	Number of BLOB frames in the tablespace.
NUM_FREE_FRAM ES	Number of free BLOB frames available in the tablespace.

SYSTEXTINDEX

The SYSTEXTINDEX table contains information on text indexes.

COLUMN NAME	DESCRIPTION
TEXTINDEX_NAME	Text index name.
TABLE_NAME	Table name.
TABLE_OWNER	Table owner.

Column Name	Description
COLUMN_ID	Column ID.
TEXT_BLOCK_SIZE	Text block size.
BASIC_BIT_LENGTH	Basic bit length.
EXT_BIT_LENGTH	Extended bit length.
CLUSTER_WIDTH	Cluster width.
NUM_TEXT_BLOCK	Number of text blocks.
AVG_TEXT_SIZE	Average text size.

SYSTRIGGER

The SYSTRIGGER table contains information on triggers.

Column Name	Description
TBNAME	Table name.
TBOWNER	Table owner.
TRIGNAME	Trigger name.
TRIGEVENT	Trigger event. 1 — Insert event 2 — Delete event 3 — Update event trigger 4 — Update column event
NUM_COL	Number of columns.
SCOL_NUM	The lowest column number updated for trigger.
TRIGTYPE	Trigger type. 1 — BEFORE and FOR EACH STATEMENT 2 — BEFORE and FOR EACH ROW 4 — AFTER and FOR EACH STATEMENT 8 — AFTER and FOR EACH ROW

COLUMN NAME	DESCRIPTION
STATUS	Status. (1: enable; 0: disable.)
OLD	Old value.
NEW	New value.
MODE	1: valid trigger; 0: invalid trigger.
TRIGDEF	Trigger definition.

SYSTRPDEST

The SYSTRPDEST table contains information on schedules used by asynchronous table replication.

COLUMN NAME	DESCRIPTION
SVRNAME	Remote database name.
USER_NAME	User account in the remote database.
STATUS	Status of the remote database. (0: normal; 1: suspend)
BEGTIME	Beginning time of replication.
INTERVAL	Interval between replicating.

SYSTRPJOB

The SYSTRPJOB table contains information on logs used by asynchronous table replication.

COLUMN NAME	DESCRIPTION
DESTINATION	The database to which data is replicated.
FN	The file numbers of a transaction log record.
OFFSET	The Offset in the transaction log record.

SYSTRPPOS

The SYSTRPPOS table contains information used by asynchronous table replication.

Column Name	Description
POSARRAY	Internal usage.

SYSTXNINFO

The SYSTXNINFO table, subset of SYSINFO table, contains information on transactions for a database.

Column Name	Description
NUM_MAX_TRANX	Maximum number of transactions.
NUM_ACT_TRANX	Number of currently active transactions.
NUM_STARTED_TRANX	Number of started transactions.
NUM_COMMITED_TRANX	Number of committed transactions.
NUM_ABORTED_TRANX	Number of aborted transactions.
NUM_COMMIT_WAITER	Number of transactions waiting commit.

SYSUSER

The SYSUSER table contains information on the status of all users currently connected to a database. Before killing a connection, you should query the SYSUSER table for the ID of the connection you want to kill. If your login host name is not registered in the network, LOGIN_HOST is anonymous.

Column Name	Description
CONNECTION_ID	Connection ID.
USER_NAME	Login user name.

COLUMN NAME	DESCRIPTION
LOGIN_TIME	Login time.
LOGIN_IP_ADDR	Login IP address.
LOGIN_HOST	Login host name.
NUM_SCAN	Number of **SELECT** operations.
NUM_INSERT	Number of **INSERT** operations.
NUM_UPDATE	Number of **UPDATE** operations.
NUM_DELETE	Number of **DELETE** operations.
NUM_TRANX	Number of processed transactions.
NUM_JBYTE_PER_TRAN	Number of journal bytes per transaction.
FIRST_W_JNR_FN	First journal file number of one active transaction.
FIRST_W_JNR_BN	First journal block number of one active transaction
NUM_BYTE_JNR_DATA	Total journal bytes used in the active transaction
NUM_J_BLOCK_DURATN	The span between first journal block number used by the active transaction and one used most recently.
SQL_CMD	The most recently executed SQL command and the command status. The command status could be the following: [PRE] - The SQL command is preparing. [EXEC] - The SQL command is executing from a SQLExecute call. [EXDIR] - The SQL command is executing from a SQLExecDirect call. [FETCH] - The operation is in a fetch data phase. [EXIT] - The SQL command has finished a prepare, execute or fetch operation.

Column Name	Description
TIME_OF_SQL_CMD	The time when the most recently used SQL command was executed.

SYSUSERFUNC

The SYSUSERFUNC table contains information on user-defined and built-in functions.

Column Name	Description
MODE	Type of function: 1 — built-in function 0 — not built-in function
FILE_NAME	The name of the file that the built-in function is in.
FUNC_NAME	The name of the built-in function.
RETURN_TYPE	Data type the built-in function returns.
NUM_OF_PARAMETER	The number of parameters in the function
PARAMETER	Data type of each parameter. The number of parameters is given by the value of NUM_OF_PARAMETER.

SYSVIEWDATA

The SYSVIEWDATA table gives information on the table views in the database.

Column Name	Description
VIEW_NAME	Table view name.
VIEW_OWNER	Table view owner.
STATUS	**0** — invalid view. **1** — valid view.

VIEW_DEFINITION	View definition.

SYSWAIT

The SYSWAIT table gives the status for locks that are currently waiting for another lock to release an object.

COLUMN NAME	DESCRIPTION
WAITING_CONNECTION	ID of connection which is waiting.
WAITED_CONNECTION	ID of connection which is being waited for.

C System Limitations

DBMaker has certain limitations to the length of names in the database, the size of indices, tables, and memory buffers, the size of files and the number of concurrent transactions. These limitations and others are summarized in the following sections.

C.1 Naming Limitations

The ANSI/ISO standard for the SQL language specifies that database objects should be given unique names, and defines the database objects that require names. These names are used in SQL statements to identify which objects the statements act on.

The database objects that require names are:

- Tables
- Columns
- Users

DBMaker conforms to the ANSI/ISO standard with the exception of user names and passwords, which may only contain 1 to 32 characters. In the ANSI/ISO standard, SQL database object names have a maximum length of 32 characters, and may contain letters and numbers. They may not contain spaces or punctuation characters.

DBMaker also supports names for several database objects and extends the range of characters that can be used:

- Indexes

- Cursors
- Tablespaces
- Primary/Foreign keys

In DBMaker, database names may contain 1 to 32 alphanumeric characters or the underscore character, in any position including the first.

All other identifiers, with the exception of passwords, may include 1 to 32 alphanumeric characters, Chinese double-byte characters, spaces, underscores, and the symbols $ and #, in any position including the first. If spaces are used, the name must be enclosed in double quotes (" "), and any trailing spaces are ignored. Passwords also follow these rules, but are limited to a maximum length of 16 characters.

The size limitations on all of the named database objects supported by DBMaker are shown in the table below.

ITEM	MINIMUM	MAXIMUM
Database Name	1	32
Tablespace Name	1	32
Table Name	1	32
Column Name	1	32
Index Name	1	32
Cursor Name	1	32
User Name	1	32
User Password	1[1]	16
Physical File Name with file path	1	256[2]
Logical File Name	1	32

[1] If a user does not set the password, then the length of the password is NULL.

[2] Including null terminator.

ITEM	MINIMUM	MAXIMUM
SQL Statement	N/A	32767

Table C-1: Minimum/maximum lengths of database object names (in characters).

C.2 Storage Limitations

The following table shows the storage limitations placed on database objects by DBMaker. While the maximum value is shown as a specific number for most of these limitations, you should keep in mind that physical system limitations, such as system memory or disk space, and operating system limitations, such as system resources or other limitations, may impose a restriction before the specified value is reached. Unless otherwise noted, all limitations are the same for all platforms supported by DBMaker.

ITEM	MINIMUM	MAXIMUM
Size of a database	—	70,000 GB
Number of files in a database	1	32767
Number of tablespaces in a database	1	32767
Number of files in a tablespace	1	32767
Number of tables in a tablespace	0	no limit[3]
Number of pages in a data file	2	524287
Number of columns in a table	1	252
Size of a tuple (row) in a table	0	3996[4]
Number of indices on a table	0	no limit[3]
Number of columns in an index	1	16
Length of the key in an index	0	1024[1]

[3] The number of tables and indexes is currently restricted only by operating system limitations.

[4] Including header. Refer to section 6.9 for more information.

ITEM	MINIMUM	MAXIMUM
Column ID which can be used in an index	1	127[5]
Number of system temporary files	1	8
Number of journal files	1	8
Number of pages in the journal file	23	524287
Number of projection columns	1	252
Number of GROUP BY columns	1	128
Number of ORDER BY columns	1	128
Number of ODBC binding parameters	0	255
Number of SQL sources	1	31[6]
Number of bytes in a BLOB file	0	2^{31}-1
Number of pages in a data buffer	15	depends on OS
Number of pages in a journal buffer	16	depends on OS
Number of operators in predicate for a scan	1	100

Table C-2: Minimum and maximum size of database objects (in bytes)

C.3 Processing Limitations

The following table shows the processing limitations placed on a database by DBMaker.

ITEM	MINIMUM	MAXIMUM
Number of concurrent transactions (connections) in	0	240

[5] Must be one of the first 127 columns.

[6] A source is defined as a database object that physically contains data. If an *SQL* statement lists an object (such as a view) that is a combination of data from other objects that actually contain the data, it should be counted as the number of objects combined to create it when calculating the number of sources.

a running database		
Length of CHAR or BINARY data items	0	3992 bytes
Length of VARCHAR data items	0	3992 bytes
Length of LONG VARCHAR or LONG VARBINARY data items	0	2^{31} - 1 bytes
Number of items in the projection list of a selected command	1	252

Index

D

E

F

I

J

K