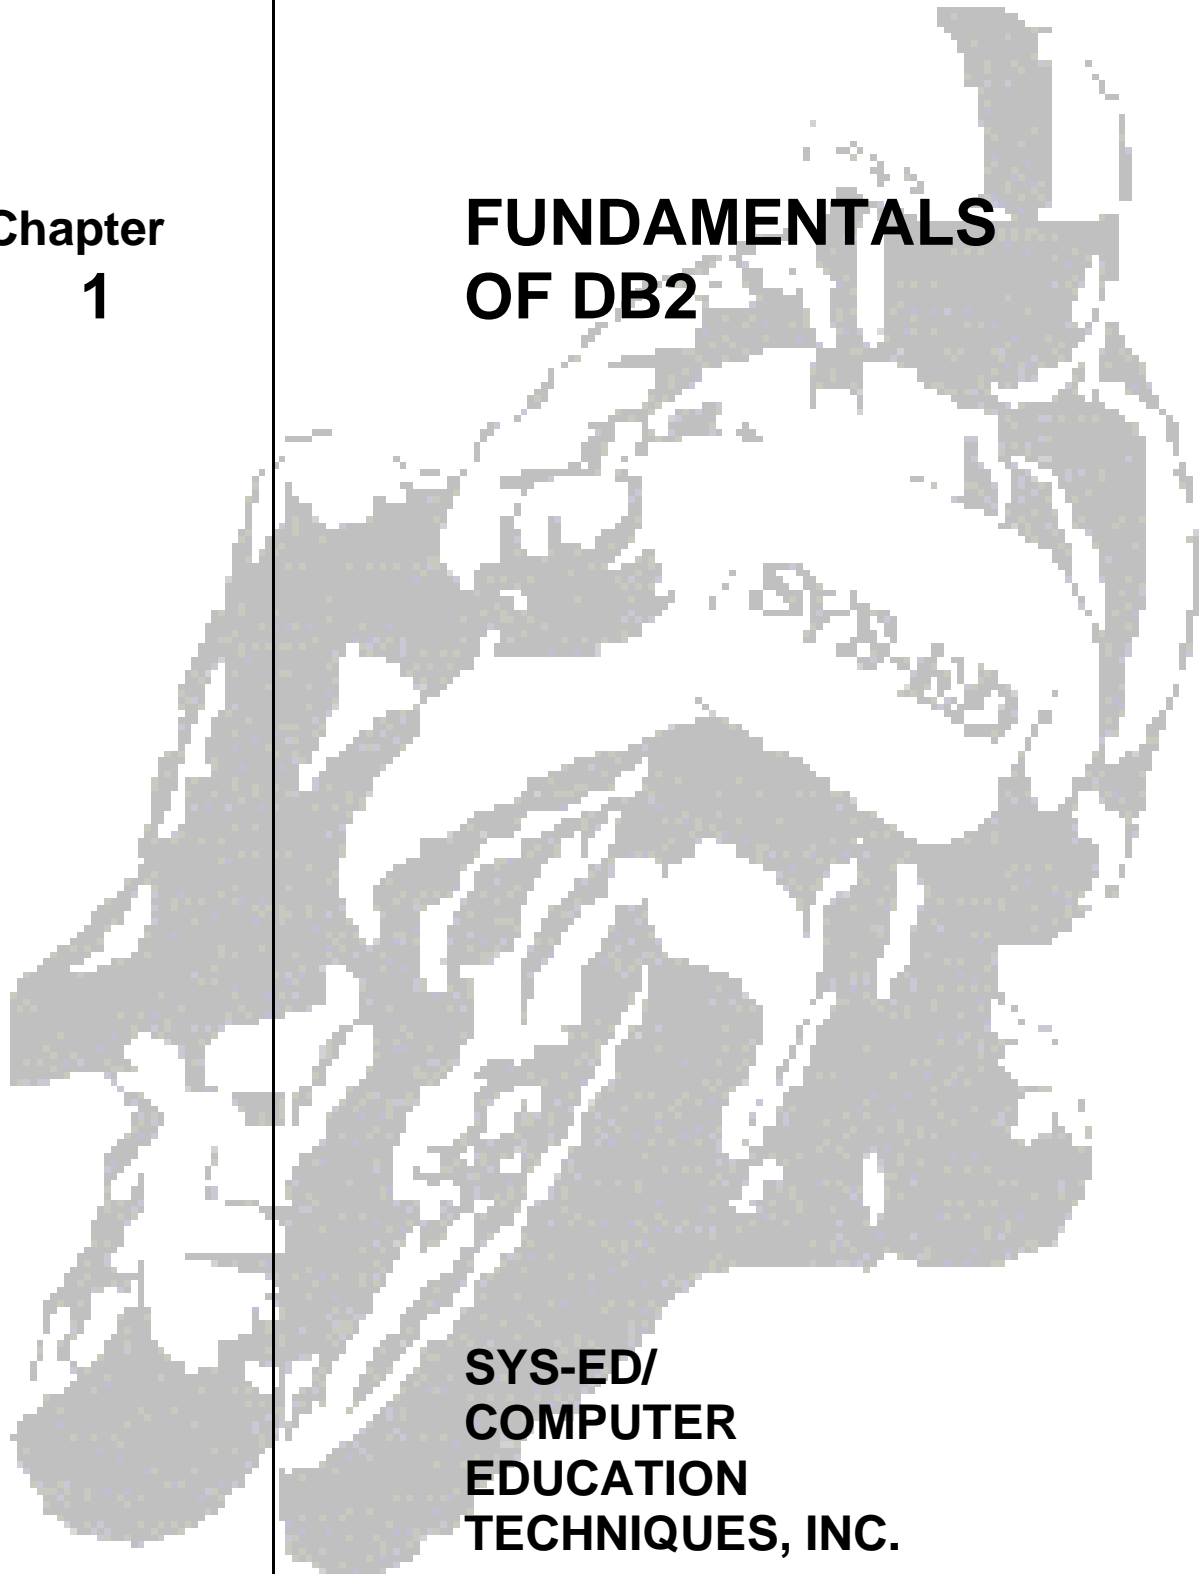


**Chapter
1**

**FUNDAMENTALS
OF DB2**



**SYS-ED/
COMPUTER
EDUCATION
TECHNIQUES, INC.**

Objectives

You will learn:

- History of DB2.
- Benefits and use of DB2.
- DB2 environment and subsystems.
- Relationship between DB2 and other systems: TSO, CICS, IMS.
- Products that work with DB2.
- Role and purpose of SQL: Structured Query Language.
- Java Database Connectivity and Embedded SQL for Java programs.
- Role and purpose of tables and base tables.
- DB2 objects: views, aliases, and indexes.
- DB2 UDB clients.
- Privileges.

1 Relational Data Base Systems

A relational data base system is a collection of tables. The mathematical term is known as a relation.

Within a table, rows are records and columns are fields.

There can be only one value at the intersection of each row and column of a table. Unlike a hierarchical data base such as IMS or IDMS, the connection between tables of a data base is through a common column (field).

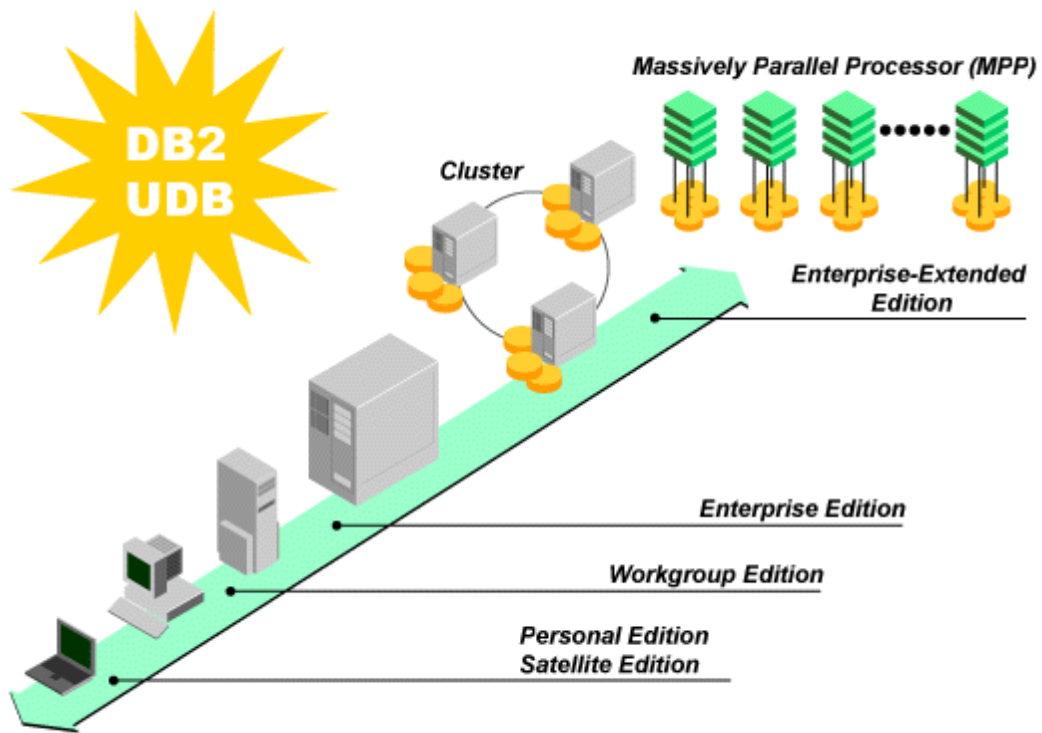
The rules underlying a relational data base system were originally stated by Dr. E. F. Codd in June 1970.

These rules have been implemented to varying degrees by the various suppliers of relational data bases.

1.1 IBM's Relational Data Base Products

System R	System R was the original relational IBM database.
SQL	This is the language which accesses the relational data base under DB2/SQL or SQL/DS.
DB2/SQL	IBM's relational database product which operates under MVS.
SQL/DS	IBM's relational database product which operates under DOS/VSE or VM.
DB2 for OS/390 and z/OS	IBM's relational database product which operates under OS/390 or z/OS.

2 DB2 at a Glance

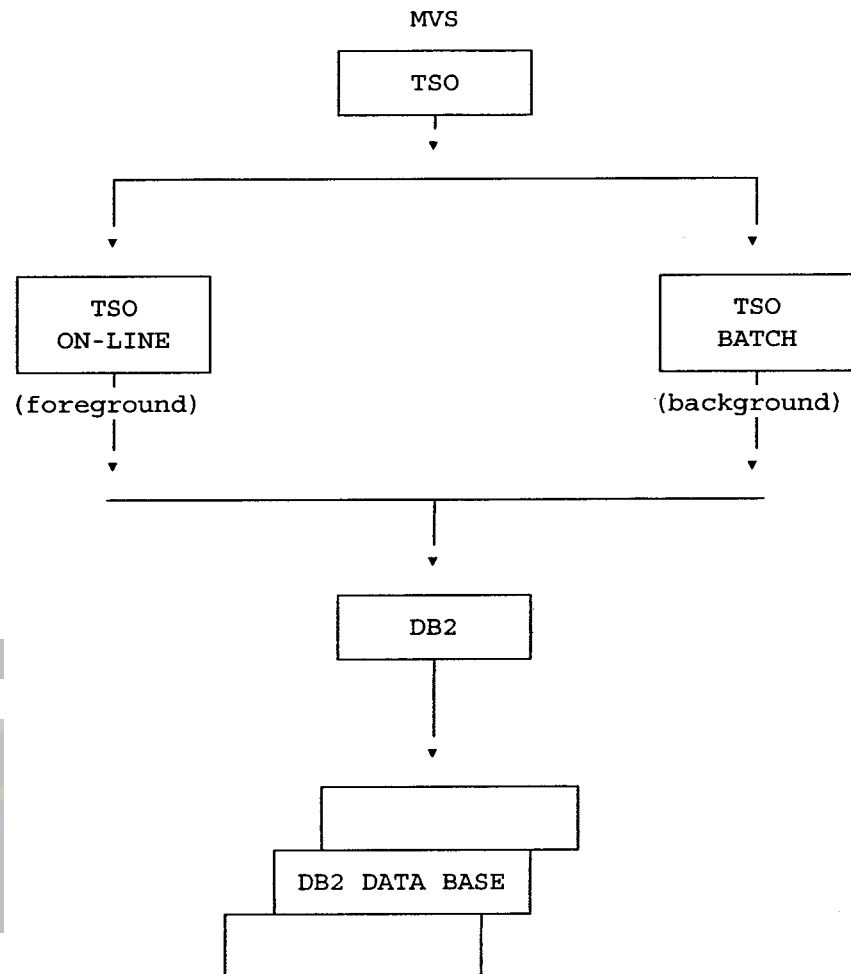


3 Major Benefits of DB2

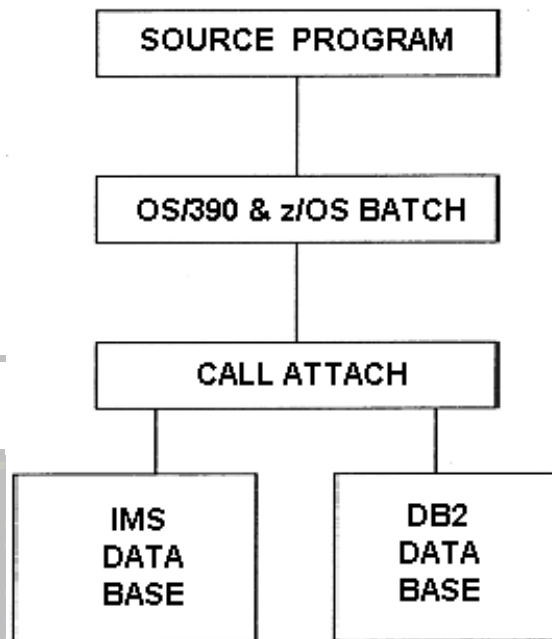
The major benefits of DB2 UDB are:

- C Runs under MVS, OS/390 and z/OS.
- C Easy to define and change.
- C High degree of data independence.
- C Easy to understand in terms of tables.
- C Provides a high level, non-procedural language for database definition and access.
- C Unlike IMS and IDMS, there are no pointers.
- C Optimizer.
- C Scrollable Cursor.
- C Limited Fetch.
- C Enhanced stored procedures.
- C JAVA support.
- C XML Extender support.

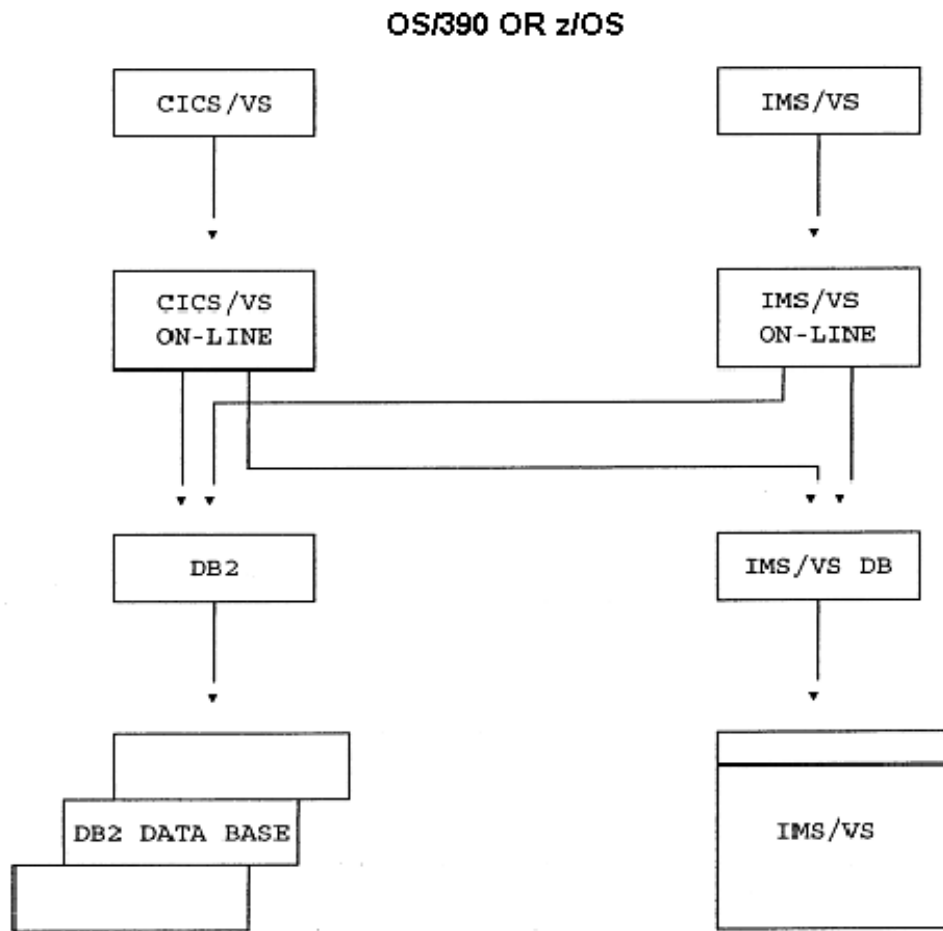
4 The DB2 Operational Environment



4.1 MVS Batch Job Submission



4.2 OS/390 or z/OS: CICS, IMS, and DB2



5 Products Which Work with DB2

The following products provide access to and work with DB2:

SQL	Structured Query Language	SQL is a non-procedural language which can be used for DB2 data base definition, data manipulation, and data control.
DB2I	Interactive DB2	Interactive DB2. DB2I is an extension to TSO which provides the capability for interactive set up and execution of SQL programs. SPUFI is used to issue and test SQL statements.
QMF	Query Management Facility	QMF is a query and report writer designed for users of DB2 databases. Operating under the security and authorization control of DB2, QMF enables updates, deletes, and insertions to be made to the database.
DXT	Data Extract	DXT makes data from IMS/DB, VSAM, or QSAM available in relational format.
CSP	Cross System Product	Fourth generation language that generates on line transaction processing systems.
DB2 Connect	DB2 Connect	DB2 Connect makes your company's host data directly available to your Personal Computer and LAN-based workstations.
PM	DB2 Performance Monitor	Runs a performance analysis on log and RMF data. This information is printed in a number of different reports.
PA	IBM DB2 SQL Performance Analyzer	Extensive analysis of SQL queries without executing them. This analysis aids you in tuning your queries to achieve maximum performance.
Utilities	DB2 Utilities Suite	All the functions of DB2 Operational Utilities and DB2 Diagnostic and Recovery Utilities in a single package.
Path Checker	IBM DB2 Path Checker	Identifies potential effects of performing a bind on one or many programs. DB2 Path Checker initiates an EXPLAIN into the plan table of the new DBRM and provides you with a report on path changes.

6 SQL: Structured Query Language

SQL is a standardized language for defining and manipulating data in a relational database. In accordance with the relational model of data, the database is perceived as a set of tables, relationships are represented by values in tables, and data is retrieved by specifying a result table that can be derived from one or more base tables.

SQL statements are executed by a database manager. One of the functions of the database manager is to transform the specification of a result table into a sequence of internal operations that optimize data retrieval. The transformation occurs in two phases: preparation and binding.

All executable SQL statements must be prepared before they can be executed. The result of preparation is the executable or operational form of the statement. The method of preparing an SQL statement and the persistence of its operational form distinguish static SQL from dynamic SQL.

7 Embedded SQL

Embedded SQL statements are SQL statements written within application programming languages such as C and COBOL and preprocessed by an SQL preprocessor before the application program is compiled. There are two types of embedded SQL: static and dynamic.

7.1 Static SQL

The source form of a static SQL statement is embedded within an application program written in a host language such as COBOL. The statement is prepared before the program is executed and the operational form of the statement persists beyond the execution of the program.

A source program containing static SQL statements must be processed by an SQL precompiler before it is compiled. The precompiler turns the SQL statements into host language comments, and generates host language statements to invoke the database manager. The syntax of the SQL statements is checked during the precompile process.

The preparation of an SQL application program includes:

1. Precompilation.
2. The binding of its static SQL statements to the target database.
3. Compilation of the modified source program.

7.2 Dynamic SQL

Programs containing embedded dynamic SQL statements must be precompiled like those containing static SQL, but unlike static SQL, the dynamic SQL statements are constructed and prepared at run time.

The SQL statement text is prepared and executed using either the PREPARE and EXECUTE statements, or the EXECUTE IMMEDIATE statement. The statement can also be executed with the cursor operations if it is a SELECT statement.

8 Java Database Connectivity and Embedded SQL for Java Programs

DB2 Universal Database implements two standards-based Java programming APIs:

JDBC	Java Database Connectivity
SQLJ	Embedded SQL for Java

Both can be used to create Java applications and applets that access DB2.

JDBC calls are translated to calls to DB2 CLI through Java native methods. JDBC requests flow from the DB2 client through DB2 CLI to the DB2 server. Static SQL cannot be used by JDBC.

SQLJ applications use JDBC as a foundation for such tasks as connecting to databases and handling SQL errors, but can also contain embedded static SQL statements in the SQLJ source files.

An SQLJ source file has to be translated with the SQLJ translator before the resulting Java source code can be compiled.

9 Tables

Tables are logical structures maintained by the database manager. Tables are made up of columns and rows. The rows are not necessarily ordered within a table (order is determined by the application program). At the intersection of every column and row is a specific data item called a value.

A column is a set of values of the same data type. A row is a sequence of values such that the *n*th value is a value of the *n*th column of the table.

9.1 Base Table

A base table is created with the CREATE TABLE statement and is used to hold persistent user data.

A result table is a set of rows that the database manager selects or generates from one or more base tables to satisfy a query.

A summary table is a table that is defined by a query that is also used to determine the data in the table. Summary tables can be used to improve the performance of queries. If the database manager determines that a portion of a query could be resolved using a summary table, the query may be rewritten by the database manager to use the summary table. This decision is based on certain settings such as CURRENT REFRESH AGE and CURRENT QUERY OPTIMIZATION special registers.

A table can have the data type of each column defined separately, or have the types for the columns based on the attributes of a user-defined structured type. This is called a typed table. A user-defined structured type may be part of a type hierarchy.

A subtype is said to inherit attributes from its supertype. Similarly, a typed table can be part of a table hierarchy. A subtable is said to inherit columns from its supertable. The term subtype applies to a user-defined structured type and all user-defined structured types that are below it in the type hierarchy.

A proper subtype of a structured type T is a structured type below T in the type hierarchy. Similarly the term subtable applies to a typed table and all typed tables that are below it in the table hierarchy. A proper subtable of a table T is a table below T in the table hierarchy.

10 Views

A view provides an alternative way of looking at the data in one or more tables.

A view is a named specification of a result table. The specification is a SELECT statement that is executed whenever the view is referenced in an SQL statement. Thus, a view can be thought of as having columns and rows just like a base table. For retrieval, all views can be used just like base tables.

Whether a view can be used in an insert, update, or delete operation depends on its definition as explained in the description of CREATE VIEW.

When the column of a view is directly derived from a column of a base table, that column inherits any constraints that apply to the column of the base table. If a view includes a foreign key of its base table, INSERT and UPDATE operations using that view are subject to the same referential constraint as the base table.

If the base table of a view is a parent table, DELETE and UPDATE operations using that view are subject to the same rules as DELETE and UPDATE operations on the base table.

A view can have the data type of each column derived from the result table, or have the types for the columns based on the attributes of a user-defined structured type. This is called a typed view. Similar to a typed table, a typed view can be part of a view hierarchy.

A subview is said to inherit columns from its superview. The term subview applies to a typed view and all typed views that are below it in the view hierarchy. A proper subview of a view V is a view below V in the typed view hierarchy.

A view may become inoperative, in which case it is no longer available for SQL statements.

11 Aliases

An alias is an alternate name for a table or view. It can be used to reference a table or view in those cases where an existing table or view can be referenced.

As with tables and views, an alias may be created, dropped, and have comments associated with it. Unlike tables, aliases may refer to each other in a process called chaining. Aliases are publicly referenced names so no special authority or privilege is required to use an alias. Access to the tables and views referred to by the alias, however, still require the appropriate authorization for the current context.

In addition to table aliases, there are other types of aliases such as database and network aliases. Aliases can also be created for nicknames.

12 Indexes

An index is an ordered set of pointers to rows of a base table. Each index is based on the values of data in one or more table columns. An index is an object that is separate from the data in the table.

When an index is created, the database manager builds this structure and maintains it automatically.

Indexes are used by the database manager to:

C Improve performance.

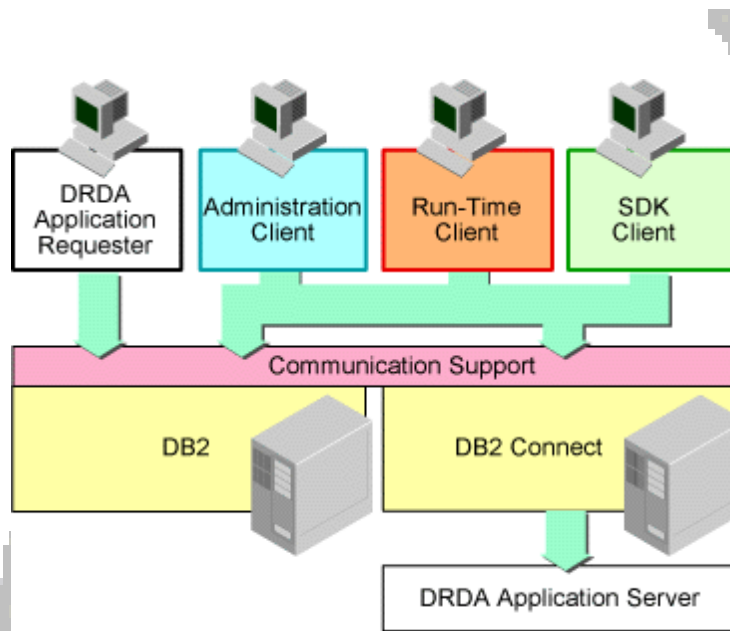
In most cases, access to data is faster than without an index.

C Ensure uniqueness.

A table with a unique index cannot have rows with identical keys.

An index cannot be created for a view. However, an index created for a table on which a view is based may improve the performance of operations on the view.

13 DB2 UDB Clients



14 Privileges

