

Introduction to Relational Databases in SQL

Candidate key: A minimal set of one or more columns whose values uniquely identify rows in a table and from which a primary key can be chosen

Cardinality: A description of how many instances of one entity can relate to instances of another (e.g., one-to-many, many-to-many)

Column (Attribute): A named vertical field in a table that defines a data element for every row, including its data type and any constraints

Composite key: A primary or candidate key composed of two or more columns whose combined values uniquely identify a row

Data type: The definition of the kind of values a column can hold (e.g., integer, text, date), which constrains storage and allowed operations

Database Management System (DBMS): Software (such as PostgreSQL) that creates, manages, and enforces rules on databases and exposes a query interface like SQL

Domain: The set of permissible values for a column as defined by its data type and any additional constraints

Independence (of events): A property where the occurrence of one event does not change the probability of another event

Explicit measure: A measure explicitly written in DAX by the user with a custom name and logic, which is reusable, unambiguous, and preferred for complex models

Foreign key: A column or set of columns in one table whose values must match existing values of a primary key in another table, used to represent relationships between tables

information_schema: A standardized, built-in set of read-only tables provided by many DBMSs that contains metadata about databases, tables, columns, and other schema objects

INSERT INTO ... SELECT DISTINCT: An SQL pattern that inserts rows into a target table by selecting unique (distinct) values from a source table to avoid inserting duplicates

NOT NULL constraint: A column-level rule that disallows NULL values, ensuring that every row must provide a concrete value for that column

Primary key: A constraint on one or more columns that uniquely identifies each row in a table and disallows NULLs and duplicates for those columns

Referential integrity: The rule that foreign key values must reference existing rows in the related table, preventing references to non-existent records

Relational database: A structured collection of data organized into tables that model real-world entities and their relationships, allowing efficient storage, querying, and enforcement of consistency rules

Row (Record): A single horizontal entry in a table representing one instance of the entity type, composed of values for each column

Surrogate key: An artificial column added to a table (often an auto-incrementing integer) used as the table's primary key when natural attributes are unsuitable

Table: A database object that stores rows of data about a single entity type and defines columns (attributes) with specific data types and constraints

UNIQUE constraint: A rule that enforces all values in a column (or a specified set of columns) are distinct across the table, preventing duplicates