

# Lesson 2: Database System Architecture

Database system architecture refers to the design and structure of a database system that defines how data is stored, accessed, and managed. It involves several components and layers that work together to ensure data integrity, security, and efficiency in data processing and retrieval.

## Three-tier Architecture

Three-tier architecture is a well-established software design pattern that divides an application into three distinct layers: the client tier, the application tier, and the database tier. This separation enhances scalability, maintainability, and flexibility by segregating different aspects of the application into their own layers.

The client tier, also known as the presentation tier, is the topmost layer in the three-tier architecture. It is responsible for interacting with the end-users and presenting the user interface. This layer handles user inputs, sending them to the application tier, and displays the processed results. Typical components in this tier include web browsers (such as Chrome, Firefox, or Safari) for web applications, standalone desktop applications, and mobile apps. The client tier's main functions are to provide a graphical or text-based interface for user interaction, handle input collection, and present data in a comprehensible format.

The application tier, also known as the logic tier or middle tier, contains the business logic of the application. It processes user requests, performs calculations, makes decisions, and manages the flow of data between the client and database tiers. Serving as the intermediary, it ensures that user interactions are properly handled and data is processed correctly. This tier typically consists of web servers, which handle HTTP requests, and application servers, which execute business logic and communicate with the database. The application tier's functions include implementing core functionalities and rules, validating and processing inputs, facilitating communication between the client and database tiers, and integrating with other services and APIs when necessary.

The database tier, or data tier, is the bottom layer and is responsible for storing, retrieving, and managing data. It ensures data integrity, security, and consistency. Components in this tier include Database Management Systems (DBMS) such as MySQL, PostgreSQL, Oracle Database, and MongoDB, along with physical or cloud-based data storage systems. The database tier's functions encompass secure

and efficient data storage, executing queries to retrieve data requested by the application tier, managing transactions to enforce data integrity, and handling data backup, recovery, and replication to prevent data loss.

Three-tier architecture offers several advantages. It enhances scalability, as each tier can be scaled independently based on demand, such as adding more web servers in the application tier to handle increased user traffic. Maintainability is improved because changes in one tier do not directly affect the other tiers, allowing for easier updates and modifications. The architecture provides flexibility by enabling the use of different technologies in each tier, supporting a more technology-agnostic approach. Security is enhanced by isolating the database tier and controlling access through the application tier, thereby protecting sensitive data. Performance is optimized by distributing the load across tiers, with the application tier handling intensive data processing while the client tier focuses on the user interface.

In conclusion, three-tier architecture is a powerful design pattern that improves the organization, scalability, and manageability of software applications. By separating the client, application, and database tiers, developers can build robust, flexible, and maintainable systems capable of efficiently handling modern computing demands.

## Components of a DBMS

A Database Management System (DBMS) is a complex software system designed to manage, store, retrieve, and manipulate data in databases. It consists of several key components, each with specific roles and responsibilities. Understanding these components is essential for comprehending how a DBMS functions and ensures data integrity, efficiency, and security. The primary components of a DBMS include the Storage Manager, Query Processor, and Transaction Manager.

### ***Storage Manager***

The Storage Manager is responsible for managing the storage of data within the database. Its primary functions include:

**File Manager:** Manages the allocation of space on disk storage and the data structures used to represent information stored on disk. It handles the creation, deletion, and modification of files that store the data.

**Buffer Manager:** Manages the data in memory buffers, acting as an intermediary between the main memory and the disk storage. It ensures that frequently accessed data is kept in memory to improve access times and overall system performance.

**Authorization and Integrity Manager:** Enforces access control policies to ensure that only authorized users can access or modify data. It also maintains data integrity by enforcing constraints such as primary keys, foreign keys, and other rules defined in the database schema.

**Disk Manager:** Handles the physical storage of data on disk drives, including data placement, storage organization, and efficient data retrieval techniques.

### ***Query Processor***

The Query Processor is responsible for interpreting and executing database queries. Its main components include:

**Query Parser:** Analyzes and parses the SQL queries submitted by users or applications. It checks the query syntax and transforms the query into an internal representation that can be processed by the query optimizer and executor.

**Query Optimizer:** Evaluates multiple query execution plans and selects the most efficient one. It uses cost-based optimization techniques, considering factors like disk I/O, CPU usage, and memory usage to determine the best execution strategy.

**Query Executor:** Executes the optimized query plan by interacting with the storage manager to retrieve and manipulate data. It carries out the actual operations like selection, projection, join, and other SQL operations.

**SQL Engine:** The core component that understands and processes SQL commands. It converts SQL queries into a series of low-level instructions that can be executed by the database engine.

### ***Transaction Manager***

The Transaction Manager ensures that database transactions are processed reliably and adhere to the ACID (Atomicity, Consistency, Isolation, Durability) properties. Its key functions include:

**Transaction Coordinator:** Manages the execution of transactions, ensuring that all parts of a transaction are completed successfully before committing the changes to the database. If any part of a transaction fails, it ensures that the entire transaction is rolled back to maintain data integrity.

**Concurrency Control Manager:** Manages concurrent access to the database by multiple transactions. It ensures that transactions are executed in a way that avoids conflicts and maintains database consistency. Techniques like locking, timestamp ordering, and multiversion concurrency control (MVCC) are used to achieve this.

**Recovery Manager:** Ensures that the database can recover from system crashes, power failures, and other unexpected events. It uses techniques such as logging and checkpointing to maintain a record of transactions and system states, allowing the database to be restored to a consistent state after a failure.

The Storage Manager, Query Processor, and Transaction Manager are integral components of a DBMS, each playing a crucial role in the overall functionality and performance of the system. The Storage Manager handles the physical storage and retrieval of data, the Query Processor interprets and executes user queries efficiently, and the Transaction Manager ensures the reliability and integrity of transactions. Together, these components enable a DBMS to provide robust, efficient, and secure data management solutions.

## System Catalog and Metadata

A system catalog, also known as a data dictionary or metadata repository, is a critical component of a Database Management System (DBMS). It serves as a central repository for metadata, which is data about the data within the database. The system catalog ensures effective data management by maintaining detailed information about database objects and their structure.

The system catalog is a collection of tables and views that the DBMS maintains to store metadata. This metadata includes information about tables, columns, indexes, constraints, views, stored procedures, functions, triggers, users, roles, and tablespaces. Each of these components plays a crucial role in the database's operation and management. Metadata in the system catalog is categorized into several types:

structural, descriptive, administrative, and statistical. Structural metadata describes the structure of database objects, such as table schemas and columns. Descriptive metadata provides information about the data itself, including descriptions and classifications. Administrative metadata includes information related to database administration, such as user permissions and storage details. Statistical metadata contains statistical information used by the query optimizer to generate efficient query execution plans, such as data distribution statistics and cardinality.

The system catalog automatically updates whenever changes are made to the database schema or when database objects are created, modified, or dropped. It plays a pivotal role in various aspects of database management. For instance, it aids in schema management by providing a centralized repository of schema-related information. The query optimizer relies on statistical metadata to generate efficient query execution plans. Metadata about constraints and relationships helps ensure data integrity and consistency across the database. Security policies and access controls are enforced using metadata about users, roles, and privileges. Additionally, the system catalog assists in routine database maintenance tasks such as backup, recovery, and performance tuning. Metadata also provides valuable context for users to understand the database's structure and content, facilitating data exploration and discovery.

Most DBMSs provide access to the system catalog through special system tables or views, which can be queried using standard SQL. For instance, the Information Schema is a standardized set of views defined by the SQL standard that provides access to metadata across different DBMSs. Proprietary views are also available, such as **SYS.ALL\_TABLES** and **SYS.ALL\_USERS** in Oracle, or **INFORMATION\_SCHEMA.TABLES** and **INFORMATION\_SCHEMA.COLUMNS** in MySQL. Example queries to access metadata from a system catalog include listing all tables in the database, getting column details for a specific table, listing all indexes on a specific table, and viewing all users and their roles.

## Logical vs. Physical Data Independence

Data independence is a fundamental concept in database management systems (DBMS) that refers to the capacity to modify a database schema at one level without affecting the schema at the next higher level. This concept is crucial for ensuring the flexibility, scalability, and maintainability of database systems. Data independence is typically categorized into two types: logical data independence and physical data independence.

## **Logical Data Independence**

Logical data independence is the ability to change the logical schema without requiring changes to the external schemas or application programs. The logical schema represents the structure of the data as seen by the database administrators, including tables, views, indexes, and relationships. Changes to the logical schema might include adding or removing tables, altering columns, modifying constraints, or creating new views.

The key advantage of logical data independence is that it allows the database structure to evolve without necessitating changes in the user applications. For example, if new fields need to be added to a table or if the relationships between tables need to be modified, these changes can be made without requiring updates to the queries and applications that interact with the database. Logical data independence ensures that users can continue to use the database seamlessly even as the underlying structure changes.

Achieving logical data independence can be challenging because it requires a clear separation between the logical structure and the application programs. Techniques such as using views, which provide a virtual table based on the result-set of a query, and employing abstraction in application development can help achieve logical data independence.

## **Physical Data Independence**

Physical data independence refers to the ability to change the physical schema without affecting the logical schema. The physical schema defines how the data is stored in the storage devices, including file structures, indexing methods, and data compression techniques. Changes to the physical schema might involve altering the storage structures, adding new indexes to improve performance, or reorganizing data to optimize access times.

The primary benefit of physical data independence is that it allows database administrators to optimize and manage the storage of data without impacting the logical structure or the applications built upon it. For instance, the DBMS might decide to move data to a different type of storage media, change the format of storage files, or introduce new indexing techniques to enhance performance. These changes can be implemented without necessitating any modifications to the logical schema or the applications that access the database.

Achieving physical data independence is generally more feasible than logical data independence because physical changes are typically handled internally by the DBMS. Advanced DBMS features, such as automatic storage management, dynamic indexing, and query optimization techniques, facilitate physical data independence by abstracting the physical storage details from the logical structure.

### **Comparison and Importance**

Both logical and physical data independence are crucial for the efficient and effective management of database systems. Logical data independence ensures that the database can adapt to changing requirements without disrupting user applications, providing flexibility and future-proofing the database against evolving business needs. Physical data independence, on the other hand, enables performance optimization and efficient storage management, ensuring that the database remains responsive and scalable as the volume of data grows.

In summary, logical data independence focuses on protecting the user applications from changes in the logical schema, whereas physical data independence deals with safeguarding the logical schema from changes in the physical storage details. Together, these forms of data independence contribute to a robust, adaptable, and efficient database management system that can meet the evolving needs of users and applications while maintaining high performance and manageability.