

Lesson 5: Storage Structures and File Systems

Understanding storage structures and file systems is essential for efficiently managing data within a Database Management System (DBMS). This lesson explores the various storage structures used by DBMSs, the organization of data on disk, and how file systems interact with database storage mechanisms. These elements play a critical role in the performance, reliability, and scalability of a database system.

Logical and Physical Storage Structures

Understanding the difference between logical and physical storage structures is fundamental for database management. These concepts determine how data is organized, stored, and accessed, impacting the performance, scalability, and reliability of a Database Management System (DBMS).

Logical Storage Structures

Logical storage structures refer to the organization of data as perceived by users and database administrators. These structures are designed to facilitate data management and provide a way to define, organize, and manipulate data without concerning the underlying physical storage details. Key logical storage structures include tables, views, schemas, indexes, and tablespaces.

Tables are the primary logical structure for storing data in a relational database. Each table is made up of rows and columns, where rows represent individual records and columns represent the attributes of those records. Tables are designed to store data in a structured format, allowing for efficient querying and manipulation.

Views are virtual tables that provide a way to present data from one or more tables in a specific format or perspective. They do not store data themselves but rather define a query that pulls data from underlying tables. Views can simplify complex queries, enhance security by restricting access to specific data, and present data in a more user-friendly manner.

Schemas are logical containers that group related database objects, such as tables, views, indexes, and procedures. Schemas help organize database objects and provide a level of abstraction that separates the logical organization of data from the physical

storage. They also play a crucial role in managing access control and security within a database.

Indexes are auxiliary structures that improve the speed of data retrieval. They are created on one or more columns of a table and provide a mechanism for quickly locating rows without scanning the entire table. Different types of indexes, such as B-tree, hash, and bitmap indexes, offer varying benefits depending on the nature of the queries and the data distribution.

Tablespaces are logical storage units that group related data files together. They allow for the logical separation of data, which can help optimize database performance by organizing data more efficiently. Tablespaces also make it easier to manage disk space allocation and perform administrative tasks, such as backups and data migrations.

Physical Storage Structures

Physical storage structures refer to the actual files and storage formats used by the DBMS to store data on disk. These structures are concerned with how data is physically written to and read from the storage media. Key physical storage structures include data files, blocks, extents, segments, and transaction logs.

Data files are the physical files on disk where the actual database data is stored. Each tablespace consists of one or more data files. The DBMS reads from and writes to these files to store and retrieve data. Proper management of data files, including their size and location, is crucial for maintaining database performance and ensuring efficient storage use.

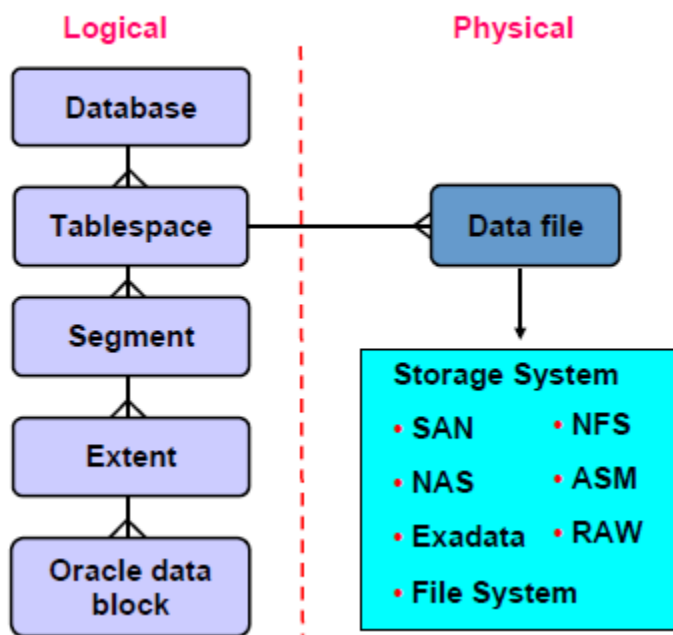
Blocks are the smallest units of storage within a data file. Also known as pages in some DBMSs, blocks typically range in size from 2KB to 32KB, depending on the database configuration. Data within a table is stored in these blocks, and the DBMS manages the reading and writing of blocks during data operations. Efficient block management is essential for optimal database performance.

Extents are larger units of storage made up of multiple contiguous blocks. An extent is allocated to a database object, such as a table or an index, to store its data. As the object grows, additional extents are allocated to accommodate new data. Managing extents efficiently helps prevent fragmentation and ensures that data is stored contiguously, which can improve read and write performance.

Segments are the highest level of physical storage structures and consist of one or more extents. Each database object, such as a table or an index, is associated with a segment. Segments group the extents belonging to a single object, facilitating efficient space management and data access.

Transaction logs are essential for maintaining the integrity and consistency of the database. They record all changes made to the database, providing a mechanism for recovery in case of system failures. Transaction logs include redo logs and undo logs. Redo logs record changes that can be reapplied during recovery, while undo logs track changes that need to be undone if a transaction is rolled back. Proper log management

involves configuring the size and location of log files and ensuring regular backups to support recovery operations.



In summary, logical and physical storage structures play a crucial role in the design, performance, and management of a DBMS. Logical storage structures, such as tables, views, schemas, indexes, and tablespaces, provide a way to organize and manage data at a higher level of abstraction. Physical storage structures, including data files, blocks, extents, segments, and transaction logs, determine how data is stored and accessed on disk.

Understanding and efficiently managing these structures are essential for ensuring a robust, high-performing, and reliable database system.

Database Files

Database files are critical components of a Database Management System (DBMS), each serving specific functions that contribute to the overall operation, integrity, and performance of the database. The primary types of database files include datafiles,

control files, and redo log files. Understanding these files and their roles is essential for effective database management.

Datafiles

Datafiles are the physical files on disk where the actual database data is stored. Each tablespace in a database is associated with one or more datafiles. These files contain all the data stored in the database, including tables, indexes, and other database objects.

Datafiles are crucial for the storage of structured data, and their efficient management directly impacts database performance. They are divided into smaller units called blocks (or pages in some DBMSs), which are the smallest units of data storage that the DBMS reads from and writes to. The size of these blocks can typically be configured based on the database's workload and performance requirements.

Proper management of datafiles involves monitoring their size and growth. As the amount of data in the database increases, datafiles need to be expanded or additional datafiles added to the tablespaces. This ensures that the database does not run out of space and can continue to function smoothly. Regular maintenance tasks, such as reorganizing data and performing backups, are essential to ensure data integrity and optimize performance.

Control Files

Control files are vital for the operation of a DBMS as they store metadata about the database. This metadata includes the structure of the database, the names and locations of datafiles and redo log files, the database creation timestamp, and the current log sequence number. Essentially, control files provide a roadmap for the DBMS, allowing it to understand and manage the database.

Control files are critical during the startup and recovery processes of the database. When a DBMS starts, it reads the control file to determine the state of the database and to locate the necessary datafiles and redo log files. In the event of a system failure, the information in the control file is used to guide the recovery process.

Given their importance, most DBMSs support multiple copies of control files stored on different disks. This redundancy ensures that if one control file becomes corrupt or is lost due to disk failure, the database can still recover using the other copies. Regular backups of control files are also recommended to protect against data loss and corruption.

Redo Log Files

Redo log files play a crucial role in maintaining the integrity and consistency of a database. They record all changes made to the database, capturing both the before and after states of modified data. This logging mechanism allows the DBMS to redo (or reapply) changes in the event of a system failure, ensuring that no committed transactions are lost.

Redo log files are organized into groups, and each group consists of one or more identical copies of redo logs (known as members). These groups are used in a circular fashion, where the DBMS writes to the current redo log file until it is full, then moves to the next file in the group. This process continues in a loop, providing continuous logging of database changes.

During a transaction, changes are first written to the redo log file before being applied to the datafiles. This ensures that even if a failure occurs before the changes are fully written to the datafiles, the DBMS can use the redo log to recover the incomplete transactions.

Proper management of redo log files includes configuring their size and number to ensure they can handle the database's transaction load. Additionally, regular archiving of filled redo log files is essential for supporting point-in-time recovery and preventing the logs from becoming a bottleneck.

Datafiles, control files, and redo log files are fundamental components of a DBMS, each playing a distinct and vital role in the system's operation. Datafiles store the actual database data, control files contain metadata essential for database management and recovery, and redo log files record all changes to ensure data integrity and support recovery processes. Effective management of these files is crucial for maintaining database performance, reliability, and data integrity. Understanding their functions and how to manage them can help database administrators ensure the smooth and efficient operation of the database system.

Tablespaces and Their Purpose

Tablespaces are an essential concept in the architecture of Database Management Systems (DBMS). They serve as logical storage units that group related data files together, enabling efficient organization, management, and optimization of data storage.

Understanding tablespaces and their purpose is crucial for effective database administration and performance tuning.

A tablespace is a logical container within a database that holds database objects such as tables, indexes, and large objects. Each tablespace consists of one or more physical datafiles, which are the actual files on disk where the database stores its data. The use of tablespaces provides a layer of abstraction between the physical storage of data and its logical organization within the database. The primary purposes of tablespaces include:

1. Logical Data Organization:

Tablespaces allow for the logical grouping of related database objects. This organization helps in managing and accessing data more efficiently. For example, a database can have separate tablespaces for different types of data such as user data, system data, and temporary data. This logical separation makes it easier to administer and optimize the database.

2. Performance Optimization:

By grouping related objects into specific tablespaces, administrators can optimize performance based on usage patterns. For instance, frequently accessed tables can be placed in tablespaces stored on high-speed storage devices, while less frequently accessed tables can be stored on slower, more cost-effective storage. This strategic allocation helps balance performance and cost.

3. Space Management:

Tablespaces facilitate efficient space management by allowing administrators to monitor and allocate disk space for different types of data. They can set quotas and manage the growth of individual tablespaces, preventing any single tablespace from consuming too much disk space and ensuring that space is available where it is most needed.

4. Backup and Recovery:

Using tablespaces enhances the flexibility and efficiency of backup and recovery operations. Administrators can back up and restore individual tablespaces rather than the entire database, which can save time and resources. This capability is particularly useful for large databases, where full backups can be time-consuming and resource-intensive.

5. Security and Access Control:

Tablespaces can also be used to enhance security and manage access control. By segregating data into different tablespaces, administrators can apply specific access

controls and security policies to each tablespace. This segregation ensures that sensitive data is stored separately and can be protected with more stringent security measures.

6. Maintenance and Administration:

Maintenance tasks such as reorganizing data, performing defragmentation, and optimizing storage can be done more efficiently at the tablespace level. Administrators can perform these tasks on individual tablespaces without affecting the entire database, leading to reduced downtime and improved overall database performance.

Practical Example of Tablespace Usage

Consider a large enterprise database that handles different types of data: transactional data, analytical data, and temporary data used for processing queries. The database administrator can create separate tablespaces for each type:

- **Transactional Tablespace:** This tablespace holds all tables and indexes related to daily transactions. It is optimized for quick read and write operations and stored on high-performance SSDs to ensure fast access.
- **Analytical Tablespace:** This tablespace contains data used for reporting and analysis. It might be stored on high-capacity HDDs that provide sufficient storage at a lower cost, as the data is accessed less frequently.
- **Temporary Tablespace:** This tablespace is used for temporary data generated during query processing, sorting operations, and intermediate result storage. It is designed for rapid data turnover and optimized for efficient space reuse.

By organizing the database in this way, the administrator can ensure that each type of data is stored optimally, improving overall database performance and manageability.

Tablespaces are a fundamental feature of DBMSs that provide logical organization and management of data. They offer numerous benefits, including improved performance, efficient space management, enhanced security, and flexible backup and recovery options. Understanding and effectively utilizing tablespaces can significantly contribute to the efficient and effective administration of a database system.

File Systems and Raw Devices in Database Management

In the context of Database Management Systems (DBMS), understanding the storage infrastructure is crucial for optimizing performance and ensuring data integrity. Two primary methods for storing database files on a disk are through file systems and raw devices. Each method has its advantages and trade-offs, and choosing the appropriate one depends on specific database requirements and workload characteristics.

A file system is a method and data structure that an operating system uses to manage files on a disk. Common file systems include NTFS (New Technology File System) on Windows, ext4 (fourth extended file system) on Linux, and APFS (Apple File System) on macOS. File systems are user-friendly and widely supported across different operating systems, allowing for straightforward file management, including naming, copying, moving, and deleting files. They are highly compatible, making installation and configuration of DBMSs simple. Additionally, file systems support various file types and sizes, offering versatility for different database needs. Modern file systems come with advanced features such as journaling, which helps in recovering from crashes by keeping track of changes not yet committed to the disk, and support for snapshots and cloning, which are useful for backups and disaster recovery.

However, file systems introduce additional overhead due to the management of file metadata, directory structures, and the handling of various file operations. This overhead can impact database performance, especially in high-throughput environments. Over time, as files are created, modified, and deleted, file systems can become fragmented, slowing down data access as the file system has to read from multiple locations on the disk.

Raw devices, also known as raw partitions, refer to disk partitions that are accessed directly by the DBMS without the intermediation of a file system. By bypassing the file system, raw devices eliminate the overhead associated with file system management, leading to significant performance improvements, particularly in high-transaction environments. Since the DBMS controls data placement directly, it can manage disk space more efficiently and reduce fragmentation, leading to more consistent performance. Moreover, without the complexity of a file system, recovery operations can be simpler and faster, as the DBMS can directly control how data is written and retrieved.

However, managing raw devices can be more complex than using file systems. It requires specialized knowledge and administrative skills to configure and maintain the storage. Raw devices are less flexible than file systems; for example, resizing raw partitions can be challenging and often requires downtime, whereas modern file

systems allow for dynamic resizing. Additionally, raw devices lack the advanced features provided by file systems, such as file permissions, encryption, and journaling. The DBMS must handle all aspects of data management, which can increase administrative overhead.

The choice between using file systems and raw devices depends on various factors, including performance requirements, administrative expertise, and specific workload characteristics. For databases with high transaction volumes and demanding performance requirements, raw devices may be the preferred choice due to their reduced overhead and direct disk access. For environments where ease of management and flexibility are more important, file systems offer significant advantages with their user-friendly features and advanced capabilities. Some organizations use a hybrid approach, employing file systems for general database storage while using raw devices for specific high-performance areas, such as redo log files or temporary storage areas.

Both file systems and raw devices have their place in database storage strategies. File systems provide ease of use, flexibility, and advanced features that simplify management and enhance security. Raw devices offer superior performance and efficiency by eliminating file system overhead. Understanding the benefits and trade-offs of each method allows database administrators to make informed decisions based on their specific needs and workload demands. Properly leveraging these storage options can optimize database performance, reliability, and scalability.