

# Lesson 7: Managing Storage Space

Monitoring storage usage in a Database Management System (DBMS) is a critical task for database administrators. It involves tracking the allocation and consumption of disk space by database objects to ensure optimal performance, prevent storage-related issues, and plan for future growth. Effective monitoring helps in making informed decisions about adding, resizing, and relocating data files, as well as managing temporary and undo tablespaces.

Monitoring storage usage is essential for several reasons. First, it helps in identifying potential storage shortages before they become critical, preventing disruptions in database operations. Second, it provides insights into the growth patterns of database objects, enabling proactive space management. Third, it aids in optimizing performance by ensuring that sufficient space is available for efficient data processing and retrieval. Finally, regular monitoring supports capacity planning, helping administrators anticipate and plan for future storage needs.

Several techniques can be employed to monitor storage usage effectively. One of the primary methods is using database views and reports. Most DBMSs provide system views and built-in reports that offer detailed information about storage usage. For example, in Oracle, views like **DBA\_DATA\_FILES**, **DBA\_FREE\_SPACE**, and **DBA\_SEGMENTS** can be queried to monitor space usage. A query like **SELECT tablespace\_name, SUM(bytes)/1024/1024 AS total\_space\_mb, SUM(bytes - NVL(free\_space, 0))/1024/1024 AS used\_space\_mb, SUM(free\_space)/1024/1024 AS free\_space\_mb FROM (SELECT tablespace\_name, bytes, bytes - NVL(free\_space, 0) AS free\_space FROM dba\_data\_files) GROUP BY tablespace\_name**; provides a summary of total, used, and free space in each tablespace.

Automated alerts and thresholds are also crucial tools for monitoring storage usage. Many DBMSs allow administrators to set up automated alerts that trigger when storage usage reaches predefined thresholds. These alerts can notify administrators of potential issues, allowing for timely intervention. Additionally, performance monitoring tools offer real-time insights into storage usage. These specialized tools provide dashboards, visualizations, and detailed reports that help in tracking and managing storage more effectively. Periodic audits and reviews of storage usage are also important, as they help maintain an up-to-date understanding of how storage is being utilized. This process can involve reviewing usage patterns, identifying underutilized or overutilized storage, and making necessary adjustments.

Temporary and undo tablespaces require particular attention in storage monitoring due to their roles in managing transient and transactional data, respectively. Temporary tablespaces store temporary data generated during operations such as sorting and hashing. Monitoring their usage involves tracking the size and usage of tempfiles, ensuring they have enough space to handle complex queries and operations without running into space issues. For example, to monitor the usage of temporary tablespaces in Oracle, a query like **SELECT tablespace\_name, SUM(blocks\*8192)/1024/1024 AS used\_space\_mb FROM v\$tempseg\_usage GROUP BY tablespace\_name;** can be used.

Undo tablespaces store undo data, which is crucial for transaction management, providing read consistency, and enabling rollback operations. Monitoring undo tablespaces involves ensuring that sufficient space is allocated to handle the undo requirements of the database workload. It is also important to monitor the undo retention settings to balance between providing adequate undo data and optimizing space usage. For example, to monitor undo tablespace usage in Oracle, a query like **SELECT tablespace\_name, SUM(bytes)/1024/1024 AS used\_space\_mb FROM dba\_undo\_extents GROUP BY tablespace\_name;** can be used.

Best practices for monitoring storage usage include setting regular monitoring intervals, utilizing automation, establishing thresholds and alerts, analyzing growth patterns, and performing regular maintenance. Establishing regular intervals for monitoring storage usage, such as daily or weekly checks, helps administrators stay informed about space utilization trends. Leveraging automated tools and scripts to regularly collect and report on storage usage reduces manual effort and ensures timely data. Defining thresholds for storage usage and setting up alerts to notify administrators when these thresholds are approached or exceeded is crucial. Regularly analyzing the growth patterns of database objects helps anticipate future storage needs and make proactive adjustments. Performing regular maintenance tasks, such as resizing datafiles, adding new storage, and reorganizing tablespaces, ensures efficient storage utilization.

In conclusion, monitoring storage usage is a vital component of managing a DBMS. By keeping a close watch on how storage is utilized, database administrators can ensure that their databases operate smoothly, efficiently, and without unexpected interruptions. Employing a combination of database views, automated alerts, performance monitoring tools, and regular audits helps maintain an optimal storage environment, supporting the overall health and performance of the database system.

## Automatic Storage Management (ASM)

Automatic Storage Management (ASM) is a feature provided by Oracle Database designed to simplify the management of database storage. ASM handles the complexity of managing storage by automating tasks related to disk and file management, which enhances performance and reduces administrative overhead. It offers a unified storage solution capable of managing both database and non-database files, providing an efficient and scalable storage infrastructure.

One of the key features of ASM is its ability to automate many tasks that traditionally require manual intervention by database administrators. ASM eliminates the need for manual file system management, such as creating and resizing datafiles, by automatically managing the storage space. This automation reduces the administrative burden and minimizes the potential for human error, making storage management more efficient and reliable.

ASM employs striping and mirroring to enhance performance and data protection. Striping distributes data evenly across all available disks, balancing the I/O load and improving performance by speeding up read and write operations. Mirroring provides data redundancy, protecting against disk failures. ASM supports different levels of mirroring, including normal redundancy with two-way mirroring and high redundancy with three-way mirroring, ensuring robust data protection.

A significant advantage of ASM is its dynamic rebalancing feature, which automatically redistributes data across disks when disks are added or removed from the ASM disk group. This rebalancing operation occurs in the background, ensuring continuous database operation without requiring downtime. Dynamic rebalancing helps maintain optimal performance and efficient storage utilization by adapting to changes in the storage configuration seamlessly.

ASM manages storage using disk groups, which are collections of disks managed as a single unit. When a disk group is created, ASM handles the allocation of data across the disks in the group. This simplifies the organization of storage resources and provides a logical structure for managing disks. Administrators can create, modify, and drop disk groups as needed, making it easy to scale storage up or down based on requirements.

Creating an ASM disk group involves specifying the name of the disk group, the redundancy level, and the disks included in the group. For example, in Oracle, the command **CREATE DISKGROUP data NORMAL REDUNDANCY DISK '/dev/sd1', '/dev/sd2', '/dev/sd3'**; creates an ASM disk group named **data** with normal redundancy, using the specified disks. Managing ASM disk groups includes tasks such as adding or

removing disks, resizing disk groups, and monitoring disk group usage. For instance, adding a new disk to an existing ASM disk group can be accomplished with the command **ALTER DISKGROUP data ADD DISK '/dev/sd4'**;, which integrates the new disk into the **data** disk group and triggers automatic rebalancing.

Monitoring ASM disk groups is crucial for ensuring optimal performance and storage utilization. Oracle provides several views and tools for this purpose. For example, querying the **V\$ASM\_DISKGROUP** view with **SELECT name, total\_mb, free\_mb, usable\_file\_mb, offline\_disks FROM v\$asm\_diskgroup;** provides information about the disk group name, total size, free space, usable file space, and the number of offline disks. This information helps administrators effectively monitor and manage ASM disk groups, ensuring that the storage infrastructure operates efficiently.

ASM offers numerous benefits that enhance the efficiency and reliability of database storage management. By automating routine tasks and providing advanced features like striping, mirroring, and dynamic rebalancing, ASM significantly reduces the complexity and administrative overhead associated with storage management. Its ability to manage both database and non-database files within a unified framework further simplifies storage operations. ASM also enhances performance through efficient data distribution and improves data protection with robust redundancy options, making it an invaluable tool for organizations seeking to streamline their storage management processes and ensure the high availability and performance of their Oracle databases.

Automatic Storage Management (ASM) is a powerful feature that simplifies and enhances the management of storage in Oracle databases. By automating complex storage tasks, providing advanced features like striping and mirroring, and ensuring dynamic rebalancing, ASM significantly reduces administrative overhead and improves performance. With its robust management capabilities and ease of use, ASM is an essential tool for database administrators looking to optimize their storage infrastructure and ensure the efficient operation of their databases.

## Reclaiming unused space (shrink, compress)

Reclaiming unused space in a Database Management System (DBMS) is essential for maintaining optimal performance and efficient use of storage resources. Over time, databases can accumulate unused space due to deleted records, outdated data, and other factors. This unused space can lead to wasted storage and degraded

performance if not managed properly. Two primary methods for reclaiming this space are shrinking and compressing database objects.

Shrinking involves reducing the size of database objects, such as tables and indexes, by reclaiming unused space. This process can improve performance by reducing the amount of space the database needs to scan during queries and other operations. Shrinking is particularly useful in environments where data is frequently deleted or updated. To shrink a table in Oracle, the **ALTER TABLE ... SHRINK SPACE** command is used, which compacts the table by moving rows closer together and adjusting the high-water mark to reflect the end of the used space in the table. For example, **ALTER TABLE example\_table SHRINK SPACE;** compacts the example\_table by reclaiming unused space. Similarly, indexes can be shrunk using the **ALTER INDEX ... SHRINK SPACE** command, such as **ALTER INDEX example\_index SHRINK SPACE;**, which compacts the specified index and frees up unnecessary space.

Compression reduces the amount of storage required by database objects by using various algorithms to compact data. Compressing data can significantly reduce storage costs and improve performance, especially for large databases with a high volume of data. Oracle provides several compression options, including Advanced Compression and Hybrid Columnar Compression, each suited for different data types and usage patterns. To compress a table in Oracle, the **ALTER TABLE ... COMPRESS** command is used. For example, **ALTER TABLE example\_table COMPRESS FOR OLTP;** enables compression optimized for Online Transaction Processing (OLTP) environments, reducing storage costs while maintaining performance. Compression can also be applied during data loading operations, such as with the **CREATE TABLE ... AS SELECT** statement. For instance, **CREATE TABLE compressed\_table COMPRESS FOR OLTP AS SELECT \* FROM example\_table;** creates a new compressed table **compressed\_table** populated with data from **example\_table**.

Regular maintenance is crucial for reclaiming unused space through shrinking and compressing database objects. Performing these operations periodically ensures that unused space is reclaimed consistently. Monitoring space usage using database views and monitoring tools helps identify objects with significant unused space. Automating shrinking and compression tasks using scheduled jobs ensures these operations are performed regularly and consistently. Before applying shrink and compression operations in a production environment, it is essential to test them in a staging environment to ensure they do not adversely affect performance. It is also important to balance the trade-offs between performance and storage savings when choosing compression options, as higher compression levels can impact performance.

In conclusion, reclaiming unused space through shrinking and compressing database objects is crucial for maintaining an efficient and high-performing database. Regularly performing these operations ensures optimal use of storage resources, reducing costs and improving overall system performance. Employing best practices such as regular maintenance, monitoring space usage, and automating tasks helps in managing storage effectively and keeping the database environment healthy.

## Partitioning strategies for large datasets

Partitioning is a crucial strategy for managing large datasets in a database. By dividing a large table into smaller, more manageable pieces called partitions, you can improve performance, simplify maintenance, and enhance the scalability of the database. Different partitioning strategies cater to various use cases and data distribution patterns. Here are some common partitioning strategies and their benefits.

### Range Partitioning

Range partitioning involves dividing a table into partitions based on a specified range of values. This method is particularly effective for time-series data or any dataset with a natural range-based key, such as dates or numerical sequences. For example, a sales table can be partitioned by month or year, making it easier to manage and query data for specific time periods.

```
CREATE TABLE sales (  
    sale_id NUMBER,  
    sale_date DATE,  
    amount NUMBER  
)  
PARTITION BY RANGE (sale_date) (  
    PARTITION sales_q1 VALUES LESS THAN (TO_DATE('2022-04-01', 'YYYY-MM-DD')),  
    PARTITION sales_q2 VALUES LESS THAN (TO_DATE('2022-07-01', 'YYYY-MM-DD')),  
    PARTITION sales_q3 VALUES LESS THAN (TO_DATE('2022-10-01', 'YYYY-MM-DD')),  
    PARTITION sales_q4 VALUES LESS THAN (TO_DATE('2023-01-01', 'YYYY-MM-DD'))  
);
```

In this example, the **sales** table is partitioned into four quarters based on the **sale\_date** column.

## Hash Partitioning

Hash partitioning distributes data evenly across partitions using a hash function. This method is suitable for scenarios where uniform data distribution is desired, and there is no natural range-based key. Hash partitioning helps in balancing the load and improving query performance by ensuring that data is evenly distributed.

```
CREATE TABLE users (  
    user_id NUMBER,  
    user_name VARCHAR2(50),  
    registration_date DATE  
)  
  
PARTITION BY HASH (user_id) (  
    PARTITION p1,  
    PARTITION p2,  
    PARTITION p3,  
    PARTITION p4  
);
```

Here, the **users** table is partitioned using the **user\_id** column, and the data is evenly distributed across four partitions.



## List Partitioning

List partitioning involves dividing a table into partitions based on a predefined list of values. This strategy is useful when you have a set of discrete values that can be grouped together logically. For example, you might partition a table based on regions or product categories.

```
CREATE TABLE employees (  
    employee_id NUMBER,  
    employee_name VARCHAR2(50),  
    department VARCHAR2(50)  
)  
PARTITION BY LIST (department) (  
    PARTITION sales VALUES ('SALES'),  
    PARTITION hr VALUES ('HR'),  
    PARTITION it VALUES ('IT'),  
    PARTITION finance VALUES ('FINANCE')  
);
```

In this example, the **employees** table is partitioned by the **department** column, with each department having its own partition.

## Composite Partitioning

Composite partitioning combines two or more partitioning methods, such as range-hash or range-list partitioning. This approach provides greater flexibility and can address more complex data distribution requirements. For instance, you might first partition data by range and then further sub-partition by hash within each range partition.

```

CREATE TABLE orders (
  order_id NUMBER,
  order_date DATE,
  customer_id NUMBER
)
PARTITION BY RANGE (order_date)
SUBPARTITION BY HASH (customer_id) (
  PARTITION p2022 VALUES LESS THAN (TO_DATE('2023-01-01', 'YYYY-MM-DD')) (
    SUBPARTITION p2022_1,
    SUBPARTITION p2022_2
  ),
  PARTITION p2023 VALUES LESS THAN (TO_DATE('2024-01-01', 'YYYY-MM-DD')) (
    SUBPARTITION p2023_1,
    SUBPARTITION p2023_2
  )
);

```

This example partitions the **orders** table by year (**order\_date**) and then further sub-partitions each year by **customer\_id** using hash partitioning.

## Interval Partitioning

Interval partitioning is an extension of range partitioning that automatically creates new partitions as data arrives. This method is ideal for handling continuous data streams without the need for manual partition management.

```

CREATE TABLE transactions (
  transaction_id NUMBER,
  transaction_date DATE,
  amount NUMBER
)
PARTITION BY RANGE (transaction_date)
INTERVAL (NUMTOYMINTERVAL(1, 'MONTH')) (
  PARTITION p_initial VALUES LESS THAN (TO_DATE('2022-01-01', 'YYYY-MM-DD'))
);

```

In this example, the **transactions** table initially contains a partition for data before January 2022. As new data arrives, Oracle automatically creates monthly partitions.

## Benefits of Partitioning

Partitioning large datasets provides several benefits, including improved query performance, easier management of large tables, enhanced scalability, and efficient use of resources. By reducing the amount of data scanned during queries, partitioning speeds up data retrieval and processing. It also simplifies tasks like backups, archiving, and purging old data, making database maintenance more manageable.

## Best Practices for Partitioning

- **Analyze Data Distribution:** Choose a partitioning strategy based on the natural distribution and access patterns of your data.
- **Monitor Performance:** Regularly monitor the performance of your partitions to ensure they continue to meet your requirements.
- **Use Composite Partitioning Wisely:** Combine multiple partitioning methods if your data and query patterns are complex.
- **Automate Management:** Use interval partitioning or other automated methods to reduce manual intervention for managing partitions.

In conclusion, partitioning is a powerful strategy for managing large datasets, offering significant benefits in terms of performance, scalability, and ease of maintenance. By selecting the appropriate partitioning method based on your data's characteristics and access patterns, you can optimize your database for better efficiency and manageability.