



# Database Programming with SQL

16-2

Indexes and Synonyms



# Objectives

This lesson covers the following objectives:

- Define an index and its use as a schema object
- Name the conditions that cause an index to be created automatically
- Create and execute a CREATE INDEX and DROP INDEX statement
- Create and execute a function-based index
- Create private and public synonyms

# Purpose

- Can you imagine going to a book store or a library and finding all the books stacked on shelves row after row in no real order?
- The only way to locate what you're looking for would be to examine every book in every row!
- Not very efficient, wouldn't you agree?
- You may be surprised to know that retrieval of data from a database is exactly like the rows of books.
- For every query, a full table scan happens.

# Purpose

- Fortunately there is a way to make finding data in a database more efficient.
- Oracle uses an index to speed up the retrieval of rows.
- In this lesson, you will learn how and when to create an index as well as how to delete an index.
- Also, in this lesson, you will learn how to create user-friendly names for database objects.
- Much like Internet Web addresses that eliminate having to know the Internet protocol address, synonyms provide a way to give an object a simpler name. (Sounds like aliases, doesn't it?)

# Indexes

- An Oracle Server index is a schema object that can speed up the retrieval of rows by using a pointer. Indexes can be created explicitly or automatically.
- If you do not have an index on the column you're selecting, then a full table scan occurs.
- An index provides direct and fast access to rows in a table.
- Its purpose is to reduce the necessity of disk I/O (input/output) by using an indexed path to locate data quickly.

# Indexes

- The index is used and maintained automatically by the Oracle Server. Once an index is created, no direct activity is required by the user.
- A ROWID is a base 64 string representation of the row address containing block identifier, row location in the block, and the database file identifier.
- Indexes use ROWID's because they are the fastest way to access any particular row.

# Indexes

- Indexes are logically and physically independent of the table they index.
- This means that they can be created or dropped at any time and have no effect on the base tables or other indexes.

COUNTRY_ID	COUNTRY_NAME	CAPITOL	REGION_ID
1	United States of America	Washington, DC	21
2	Canada	Ottawa	21
3	Republic of Kazakhstan	Astana	143
7	Russian Federation	Moscow	151
12	Coral Sea Islands Territory	-	9
13	Cook Islands	Avarua	9
15	Europa Island	-	18
20	Arab Republic of Egypt	Cairo	15
...	...	...	...



# Indexes

- Note: When you drop a table, corresponding indexes are also dropped.

COUNTRY_ID	COUNTRY_NAME	CAPITOL	REGION_ID
1	United States of America	Washington, DC	21
2	Canada	Ottawa	21
3	Republic of Kazakhstan	Astana	143
7	Russian Federation	Moscow	151
12	Coral Sea Islands Territory	-	9
13	Cook Islands	Avarua	9
15	Europa Island	-	18
20	Arab Republic of Egypt	Cairo	15
...	...	...	...

# Types of Indexes

- Two types of indexes can be created:
  - Unique index: The Oracle Server automatically creates this index when you define a column in a table to have a PRIMARY KEY or a UNIQUE KEY constraint.
    - The name of the index is the name given to the constraint.
    - Although you can manually create a unique index, it is recommended that you create a unique constraint in the table, which implicitly creates a unique index.
  - Nonunique index: This is an index that a user can create to speed up access to the rows.
    - For example, to optimize joins, you can create an index on the FOREIGN KEY column, which speeds up the search to match rows to the PRIMARY KEY column.

# Creating an Index

- Create an index on one or more columns by issuing the CREATE INDEX statement:

```
CREATE INDEX index_name  
ON table_name( column...,column)
```

- To create an index in your schema, you must have the CREATE TABLE privilege.
- To create an index in any schema, you need the CREATE ANY INDEX privilege or the CREATE TABLE privilege on the table on which you are creating the index. Null values are not included in the index.

# Creating an Index

- For example, to improve the speed of query access to the REGION\_ID column in the WF\_COUNTRIES table:

```
CREATE INDEX wf_cont_reg_id_idx  
ON wf_countries(region_id);
```

# When to Create an Index

- An index should be created only if:
  - The column contains a wide range of values
  - A column contains a large number of null values
  - One or more columns are frequently used together in a WHERE clause or a join condition
  - The table is large and most queries are expected to retrieve less than 2-4% of the rows.

# When Not to Create an Index

- When deciding whether or not to create an index, more is not always better.
- Each DML operation (INSERT, UPDATE, DELETE) that is performed on a table with indexes means that the indexes must be updated.
- The more indexes you have associated with a table, the more effort it takes to update all the indexes after the DML operation.

# When Not to Create an Index

- It is usually not worth creating an index if:
  - The table is small
  - The columns are not often used as a condition in the query
  - Most queries are expected to retrieve more than 2-4 % of the rows in the table
  - The table is updated frequently
  - The indexed columns are referenced as part of an expression

# Composite Index

- A composite index (also called a "concatenated" index) is an index that you create on multiple columns in a table.
- Columns in a composite index can appear in any order and need not be adjacent in the table.
- Composite indexes can speed retrieval of data for SELECT statements in which the WHERE clause references all or the leading portion of the columns in the composite index.

```
CREATE INDEX emps_name_idx  
ON employees(first_name, last_name);
```



# Composite Index

- Null values are not included in the composite index.
- To optimize joins, you can create an index on the FOREIGN KEY column, which speeds up the search to match rows to the PRIMARY KEY column.
- The optimizer does not use an index if the WHERE clause contains the IS NULL expression.

# Confirming Indexes

- Confirm the existence of indexes from the `USER_INDEXES` data dictionary view.
- You can also check the columns involved in an index by querying the `USER_IND_COLUMNS` view.
- The query shown on the next slide is a join between the `USER_INDEXES` table (names of the indexes and their uniqueness) and the `USER_IND_COLUMNS` (names of the indexes, table names, and column names) table.

# Confirming Indexes

```
SELECT DISTINCT ic.index_name, ic.column_name,  
                ic.column_position, id.uniqueness  
FROM user_indexes id, user_ind_columns ic  
WHERE id.table_name = ic.table_name  
AND ic.table_name = 'EMPLOYEES';
```

INDEX NAME	COLUMN NAME	COL POS	UNIQUENESS
EMP_EMAIL_UK	EMAIL	1	UNIQUE
EMP_EMP_ID_PK	EMPLOYEE_ID	1	UNIQUE
EMP_DEPARTMENT_IX	DEPARTMENT_ID	1	NONUNIQUE
EMP_JOB_IX	JOB_ID	1	NONUNIQUE
EMP_MANAGER_IX	MANAGER_ID	1	NONUNIQUE
EMP_NAME_IX	LAST_NAME	1	NONUNIQUE
EMP_NAME_IX	FIRST_NAME	2	NONUNIQUE

# Function-based Indexes

- A function-based index stores the indexed values and uses the index based on a SELECT statement to retrieve the data.
- A function-based index is an index based on expressions.
- The index expression is built from table columns, constants, SQL functions, and user-defined functions.

# Function-based Indexes

- Function-based indexes are useful when you don't know in what case the data was stored in the database.
- For example, you can create a function-based index that can be used with a SELECT statement using UPPER in the WHERE clause.
- The index will be used in this search.

```
CREATE INDEX upper_last_name_idx  
ON employees (UPPER(last_name));
```

```
SELECT *  
FROM employees  
WHERE UPPER(last_name) = 'KING';
```

# Function-based Indexes

- Function-based indexes defined with the UPPER(column\_name) or LOWER(column\_name) keywords allow case-insensitive searches.
- If you don't know how the employee last names were entered into the database, you could still use the index by entering uppercase in the SELECT statement.
- When a query is modified using an expression in the WHERE clause, the index won't use it unless you create a function-based index to match the expression.

# Function-based Indexes

- For example, the following statement allows for case-insensitive searches using the index:

```
CREATE INDEX upper_last_name_idx  
ON employees (UPPER(last_name));
```

```
SELECT *  
FROM employees  
WHERE UPPER(last_name) LIKE 'KIN%';
```

# Function-based Indexes

- To ensure that the Oracle Server uses the index rather than performing a full table scan, be sure that the value of the function is not null in subsequent queries.
- For example, the following statement is guaranteed to use the index, but without the WHERE clause the Oracle Server may perform a full table scan:

```
SELECT *  
FROM employees  
WHERE UPPER (last_name) IS NOT NULL  
ORDER BY UPPER (last_name);
```

- The Oracle Server treats indexes with columns marked DESC as function-based indexes.
- The columns marked DESC are sorted in descending order.



# Function-based Indexes

- All of these examples use the UPPER and LOWER functions, but it is worth noticing that while these two are very frequently used in Function Based Indexes, the Oracle database is not limited to them.
- Any valid Oracle Built-in function, for example TO\_CHAR can be used.

# Function-based Indexes

- Another example of Function Based Indexes is shown here.
- The employees table is queried to find any employees hired in 1987.

```
SELECT first_name, last_name, hire_date
FROM employees
WHERE TO_CHAR(hire_date, 'yyyy') = '1987'
```

- This query results in a Full Table Scan, which can be a very expensive operation if the table is big.
- Even if the hire\_date column is indexed, the index is not used due to the TO\_CHAR expression.

# Function-based Indexes

- Once we create the following Function Based Index, we can run the same query, but this time avoid the expensive Full Table Scan.

```
CREATE INDEX emp_hire_year_idx  
ON employees (TO_CHAR(hire_date, 'yyyy'));
```

```
SELECT first_name, last_name, hire_date  
FROM employees  
WHERE TO_CHAR(hire_date, 'yyyy') = '1987'
```

FIRST_NAME	LAST_NAME	HIRE_DATE
Steven	King	17/Jun/1987
Jennifer	Whalen	17/Sep/1987

- Now, Oracle can use the index on the hire\_date column.

# Removing an Index

- You cannot modify indexes.
- To change an index, you must drop it and then re-create it.
- Remove an index definition from the data dictionary by issuing the DROP INDEX statement.

# Removing an Index

- To drop an index, you must be the owner of the index or have the DROP ANY INDEX privilege.
- If you drop a table, indexes and constraints are automatically dropped, but views and sequences remain.

```
DROP INDEX upper_last_name_idx;
```

```
DROP INDEX emps_name_idx;
```

```
DROP INDEX emp_hire_year_idx;
```

# SYNONYM

- In SQL, as in language, a synonym is a word or expression that is an accepted substitute for another word.
- Synonyms are used to simplify access to objects by creating another name for the object.
- Synonyms can make referring to a table owned by another user easier and shorten lengthy object names.
- For example, to refer to the `amy_copy_employees` table in your classmate's schema, you can prefix the table name with the name of the user who created it followed by a period and then the table name, as in `USMA_SBHS_SQL01_S04. amy_copy_employees`.

# SYNONYM

- Creating a synonym eliminates the need to qualify the object name with the schema and provides you with an alternative name for a table, view, sequence, procedure, or other object.
- This method can be especially useful with lengthy object names, such as views.
- The database administrator can create a public synonym accessible to all users and can specifically grant the CREATE PUBLIC SYNONYM privilege to any user, and that user can create public synonyms.

# SYNONYM

- Syntax:

```
CREATE [PUBLIC] SYNONYM synonym  
FOR object;
```

- PUBLIC: creates a synonym accessible to all users
- synonym: is the name of the synonym to be created
- object: identifies the object for which the synonym is created

```
CREATE SYNONYM amy_emps  
FOR amy_copy_employees;
```



# SYNONYM Guidelines

- Guidelines:
  - The object cannot be contained in a package.
  - A private synonym name must be distinct from all other objects owned by the same user.
- To remove a synonym:

```
DROP [PUBLIC] SYNONYM name_of_synonym
```

```
DROP SYNONYM amy_emps;
```

# Confirming a SYNONYM

- The existence of synonyms can be confirmed by querying the USER\_SYNONYMS data dictionary view.

Column Name	Contents
Synonym_name	Name of the synonym.
Table_name	Owner of the object referenced by the synonym.
Table_owner	Name of the object referenced by the synonym.
Db_link	Database link referenced in a remote synonym.

# Terminology

Key terms used in this lesson included:

- Composite index
- Confirming index
- CREATE PUBLIC SYNONYM
- DROP INDEX
- Function-based index
- Non-unique index
- Synonym
- Unique index

# Summary

In this lesson, you should have learned how to:

- Define an index and its use as a schema object
- Name the conditions that cause an index to be created automatically
- Create and execute a CREATE INDEX and DROP INDEX statement
- Create function-based indexes
- Create private and public synonyms

