DML SQL Statements

Database Design

Department of Computer Engineering

Sharif University of Technology

Maryam Ramezani maryam.ramezani@sharif.edu

Projection

Table 1

Selection

Table 1

Table 1

Join

Table 2

Introduction

```sql
CREATE TABLE AGENTS (
AGENT_CODE CHAR(6) PRIMARY KEY,
AGENT_NAME CHAR(40),
WORKING_AREA CHAR(35),
COMMISSION NUMERIC(10,2),
PHONE_NO CHAR(15),
COUNTRY VARCHAR(25)
);
CREATE TABLE ORDERS (
ORD_NUM SERIAL PRIMARY KEY,
ORD_AMOUNT NUMERIC(12,2) NOT NULL,
ADVANCE_AMOUNT NUMERIC(12,2) NOT NULL,
ORD_DATE DATE NOT NULL,
CUST_CODE VARCHAR(6) NOT NULL REFERENCES
CUSTOMER,
AGENT_CODE CHAR(6) NOT NULL REFERENCES
AGENTS,
ORD_DESCRIPTION VARCHAR(60) NOT NULL
);
```

```sql
CREATE TABLE CUSTOMER (
CUST_CODE VARCHAR(6) PRIMARY KEY,
CUST_NAME VARCHAR(40) NOT NULL,
CUST_CITY CHAR(35),
WORKING_AREA VARCHAR(35) NOT NULL,
CUST_COUNTRY VARCHAR(20) NOT NULL,
GRADE INTEGER,
OPENING_AMT NUMERIC(12,2) NOT NULL,
RECEIVE_AMT NUMERIC(12,2) NOT NULL,
PAYMENT_AMT NUMERIC(12,2) NOT NULL,
OUTSTANDING_AMT NUMERIC(12,2) NOT NULL,
PHONE_NO VARCHAR(17) NOT NULL,
AGENT_CODE CHAR(6) NOT NULL REFERENCES
AGENTS
);
```

SELECT

Maryam Ramezani

Basic `SELECT` Statement

```
SELECT          *|{[DISTINCT] column|expression [alias],...}
FROM            table;
```

❑ `SELECT` identifies *what* columns

❑ `FROM` identifies *which* table

Selecting All Columns

```
SELECT *
FROM    departments;
```

DEPARTMENT_ID	DEPARTMENT_NAME	MANAGER_ID	LOCATION_ID
10	Administration	200	1700
20	Marketing	201	1800
50	Shipping	124	1500
60	IT	103	1400
80	Sales	149	2500
90	Executive	100	1700
110	Accounting	205	1700
190	Contracting		1700

8 rows selected.

Selecting Specific Columns

```
SELECT  department_id, location_id
FROM    departments;
```

DEPARTMENT_ID	LOCATION_ID
10	1700
20	1800
50	1500
60	1400
80	2500
90	1700
110	1700
190	1700

8 rows selected.

Arithmetic Expressions

Create expressions with number and date data by using arithmetic operators.

Operator	Description
+	Add
-	Subtract
*	Multiply
/	Divide

Using Arithmetic Operators

```
SELECT last_name, salary, salary + 300
FROM   employees;
```

LAST_NAME	SALARY	SALARY+300
King	24000	24300
Kochhar	17000	17300
De Haan	17000	17300
Hunold	9000	9300
Ernst	6000	6300

...

LAST_NAME	SALARY	SALARY+300
Hartstein	13000	13300
Fay	6000	6300
Higgins	12000	12300
Gietz	8300	8600

20 rows selected.

Operator Precedence & Using Parentheses

```
SELECT last_name, salary, 12*(salary+100)
FROM    employees;
```

LAST_NAME	SALARY	12*(SALARY+100)
King	24000	289200
Kochhar	17000	205200
De Haan	17000	205200
Hunold	9000	109200
Ernst	6000	73200

...

Hartstein	13000	157200
Fay	6000	73200
Higgins	12000	145200
Gietz	8300	100800

20 rows selected.

```
SELECT last_name, salary, 12*salary+100
FROM    employees;
```

LAST_NAME	SALARY	12*SALARY+100
King	24000	288100
Kochhar	17000	204100
De Haan	17000	204100
Hunold	9000	108100
Ernst	6000	72100

...

Hartstein	13000	156100
Fay	6000	72100
Higgins	12000	144100
Gietz	8300	99700

20 rows selected.

Defining a Null Value

❑ A null is a value that is unavailable, unassigned, unknown, or inapplicable.

❑ A null is not the same as zero or a blank space.

❑ Arithmetic expressions containing a null value evaluate to null.

```
SELECT last_name, 12*salary*commission_pct
FROM    employees;
```

LAST_NAME	JOB_ID	SALARY	COMMISSION_PCT
King	AD_PRES	24000	
Kochhar	AD_VP	17000	

...

LAST_NAME	JOB_ID	SALARY	COMMISSION_PCT
Zlotkey	SA_MAN	10500	.2
Abel	SA_REP	11000	.3
Taylor	SA_REP	8600	.2

...

LAST_NAME	JOB_ID	SALARY	COMMISSION_PCT
Gietz	AC_ACCOUNT	8300	

20 rows selected.

Using Column Aliases

A column alias:

❑ Renames a column heading

❑ Is useful with calculations

❑ Immediately follows the column name – there can also be the optional `AS` keyword between the column name and alias

❑ Requires double quotation marks if it contains spaces or special characters or is case sensitive

```
SELECT last_name AS name, commission_pct comm
FROM    employees;
```

NAME	COMM
King	
Kochhar	
De Haan	

...

20 rows selected.

```
SELECT last_name "Name", salary*12 "Annual Salary"
FROM    employees;
```

Name	Annual Salary
King	288000
Kochhar	204000
De Haan	204000

...

20 rows selected.

Concatenation Operator

A concatenation operator:

❑ Concatenates columns or character strings to other columns

❑ Is represented by two vertical bars (||)

❑ Creates a resultant column that is a character expression

```
SELECT      last_name||job_id AS "Employees"
FROM        employees;
```

Employees
KingAD_PRES
KochharAD_VP
De HaanAD_VP
HunoldIT_PROG
ErnstIT_PROG
LorentzIT_PROG
MourgosST_MAN
RajsST_CLERK

...

20 rows selected.

Literal Character Strings

❑ A literal is a character, a number, or a date included in the `SELECT` list.

❑ Date and character literal values must be enclosed within single quotation marks.

❑ Each character string is output once for each row returned.

```
SELECT last_name           ||' is a '||job_id
       AS "Employee Details"
FROM   employees;
```

Employee Details
King is a AD_PRES
Kochhar is a AD_VP
De Haan is a AD_VP
Hunold is a IT_PROG
Ernst is a IT_PROG
Lorentz is a IT_PROG
Mourgos is a ST_MAN
Rajs is a ST_CLERK

...

20 rows selected.

Duplicate Rows

❑ The default display of queries is all rows, including duplicate rows.

```
SELECT department_id
FROM   employees;
```

DEPARTMENT_ID
90
90
90
60
60
60
50
50
50

...

20 rows selected.

Eliminating Duplicate Rows

❑ Eliminate duplicate rows by using the DISTINCT keyword in the SELECT clause.

```
SELECT DISTINCT department_id
FROM    employees;
```

DEPARTMENT_ID
10
20
50
60
80
90
110

8 rows selected.

Restricting and Sorting

Limiting Rows Using a Selection

EMPLOYEES

EMPLOYEE_ID	LAST_NAME	JOB_ID	DEPARTMENT_ID
100	King	AD_PRES	90
101	Kochhar	AD_VP	90
102	De Haan	AD_VP	90
103	Hunold	IT_PROG	60
104	Ernst	IT_PROG	60
107	Lorentz	IT_PROG	60
124	Mourgos	ST_MAN	50

...

20 rows selected.

"retrieve all employees in department 90"

EMPLOYEE_ID	LAST_NAME	JOB_ID	DEPARTMENT_ID
100	King	AD_PRES	90
101	Kochhar	AD_VP	90
102	De Haan	AD_VP	90

Limiting the Rows Selected

❑ Restrict the rows returned by using the `WHERE` clause.

❑ The `WHERE` clause follows the `FROM` clause.

```
SELECT        *|{[DISTINCT] column|expression [alias],...}
FROM          table
[WHERE        condition(s)];
```

Using the `WHERE` Clause

❑ With fixed value

```
SELECT  employee_id, last_name, job_id, department_id
FROM    employees
WHERE   department_id = 90 ;
```

EMPLOYEE_ID	LAST_NAME	JOB_ID	DEPARTMENT_ID
100	King	AD_PRES	90
101	Kochhar	AD_VP	90
102	De Haan	AD_VP	90

❑ With variable

```
SELECT  employee_id, last_name, job_id, department_id
FROM    employees
WHERE   department_id = :input ;
```

Character Strings and Dates

❑ Character strings and date values are enclosed in single quotation marks.

❑ Character values are case sensitive, and date values are format sensitive.

❑ The default date format is DD-MON-RR.

```
SELECT  last_name, job_id, department_id
FROM    employees
WHERE   last_name = 'Whalen';
```

Comparison Conditions

Operator	Meaning
=	Equal to
>	Greater than
>=	Greater than or equal to
<	Less than
<=	Less than or equal to
<>	Not equal to

Operator	Meaning
BETWEEN ...AND...	Between two values (inclusive),
IN(set)	Match any of a list of values
LIKE	Match a character pattern
IS NULL	Is a null value

Using Comparison Conditions

```
SELECT last_name, salary
FROM   employees
WHERE  salary <= 3000;
```

LAST_NAME	SALARY
Matos	2600
Vargas	2500

❏ Use the `BETWEEN` condition to display rows based on a range of values.

```
SELECT last_name, salary
FROM   employees
WHERE  salary BETWEEN 2500 AND 3500;
```

Lower limit   Upper limit

LAST_NAME	SALARY
Rajs	3500
Davies	3100
Matos	2600
Vargas	2500

Using the `IN` Condition

❑ Use the `IN` membership condition to test for values in a list.

```
SELECT  employee_id, last_name, salary, manager_id
FROM    employees
WHERE   manager_id IN (100, 101, 201);
```

EMPLOYEE_ID	LAST_NAME	SALARY	MANAGER_ID
202	Fay	6000	201
200	Whalen	4400	101
205	Higgins	12000	101
101	Kochhar	17000	100
102	De Haan	17000	100
124	Mourgos	5800	100
149	Zlotkey	10500	100
201	Hartstein	13000	100

8 rows selected.

Using the `LIKE` Condition

❑ Use the `LIKE` condition to perform wildcard searches of valid search string values.

❑ Search conditions can contain either literal characters or numbers:
  ○ `%` denotes zero or many characters.
  ○ `_` denotes one character.

```
SELECT      first_name
FROM        employees
WHERE       first_name LIKE 'S%';
```

Using the `LIKE` Condition

❑ You can combine pattern-matching characters.

```
SELECT  last_name
FROM    employees
WHERE   last_name LIKE '_o%';
```

LAST_NAME
Kochhar
Lorentz
Mourgos

❑ You can use the `ESCAPE` identifier to search for the actual % and _ symbols.

Using the `NULL` Conditions

❑ Test for nulls with the `IS NULL` operator.

```
SELECT last_name, manager_id
FROM    employees
WHERE   manager_id IS NULL;
```

LAST_NAME	MANAGER_ID
King	

Logical Conditions

Operator	Meaning
AND	Returns TRUE if *both* component conditions are true
OR	Returns TRUE if *either* component condition is true
NOT	Returns TRUE if the following condition is false

Using the `AND` Operator

❏ `AND` requires both conditions to be true.

```
SELECT  employee_id, last_name, job_id, salary
FROM    employees
WHERE   salary >=10000
AND     job_id LIKE '%MAN%';
```

EMPLOYEE_ID	LAST_NAME	JOB_ID	SALARY
149	Zlotkey	SA_MAN	10500
201	Hartstein	MK_MAN	13000

Using the OR Operator

❑ OR requires either condition to be true.

```
SELECT  employee_id, last_name, job_id, salary
FROM    employees
WHERE   salary >= 10000
OR      job_id LIKE '%MAN%';
```

EMPLOYEE_ID	LAST_NAME	JOB_ID	SALARY
100	King	AD_PRES	24000
101	Kochhar	AD_VP	17000
102	De Haan	AD_VP	17000
124	Mourgos	ST_MAN	5800
149	Zlotkey	SA_MAN	10500
174	Abel	SA_REP	11000
201	Hartstein	MK_MAN	13000
205	Higgins	AC_MGR	12000

8 rows selected.

Using the NOT Operator

```
SELECT last_name, job_id
FROM    employees
WHERE   job_id
        NOT IN ('IT_PROG', 'ST_CLERK', 'SA_REP');
```

LAST_NAME	JOB_ID
King	AD_PRES
Kochhar	AD_VP
De Haan	AD_VP
Mourgos	ST_MAN
Zlotkey	SA_MAN
Whalen	AD_ASST
Hartstein	MK_MAN
Fay	MK_REP
Higgins	AC_MGR
Gietz	AC_ACCOUNT

10 rows selected.

Rules of Precedence

Order Evaluated	Operator
1	**Arithmetic operators**
2	**Concatenation operator**
3	**Comparison conditions**
4	`IS [NOT] NULL, LIKE, [NOT] IN`
5	`[NOT] BETWEEN`
6	`NOT` **logical condition**
7	`AND` **logical condition**
8	`OR` **logical condition**

Override rules of precedence by using parentheses.

Rules of Precedence

```
SELECT last_name, job_id, salary
FROM    employees
WHERE   job_id = 'SA_REP'
OR      job_id = 'AD_PRES'
AND     salary > 15000;
```

LAST_NAME	JOB_ID	SALARY
King	AD_PRES	24000
Abel	SA_REP	11000
Taylor	SA_REP	8600
Grant	SA_REP	7000

Rules of Precedence

❑ Use parentheses to force priority.

```
SELECT last_name, job_id, salary
FROM    employees
WHERE   (job_id = 'SA_REP'
OR       job_id = 'AD_PRES')
AND      salary > 15000;
```

LAST_NAME	JOB_ID	SALARY
King	AD_PRES	24000

ORDER BY Clause

❑ Sort rows with the `ORDER BY` clause
   ○ ASC: ascending order, default
   ○ DESC: descending order
❑ The `ORDER BY` clause comes last in the `SELECT` statement.

```
SELECT    last_name, job_id, department_id, hire_date
FROM      employees
ORDER BY hire_date ;
```

LAST_NAME	JOB_ID	DEPARTMENT_ID	HIRE_DATE
King	AD_PRES	90	17-JUN-87
Whalen	AD_ASST	10	17-SEP-87
Kochhar	AD_VP	90	21-SEP-89
Hunold	IT_PROG	60	03-JAN-90
Ernst	IT_PROG	60	21-MAY-91

...

20 rows selected.

Sorting in Descending Order

```
SELECT    last_name, job_id, department_id, hire_date
FROM      employees
ORDER BY hire_date DESC ;
```

LAST_NAME	JOB_ID	DEPARTMENT_ID	HIRE_DATE
Zlotkey	SA_MAN	80	29-JAN-00
Mourgos	ST_MAN	50	16-NOV-99
Grant	SA_REP		24-MAY-99
Lorentz	IT_PROG	60	07-FEB-99
Vargas	ST_CLERK	50	09-JUL-98
Taylor	SA_REP	80	24-MAR-98
Matos	ST_CLERK	50	15-MAR-98
Fay	MK_REP	20	17-AUG-97
Davies	ST_CLERK	50	29-JAN-97

...

20 rows selected.

Sorting by Column Alias

```
SELECT employee_id, last_name, salary*12 annsal
FROM    employees
ORDER BY annsal;
```

EMPLOYEE_ID	LAST_NAME	ANNSAL
144	Vargas	30000
143	Matos	31200
142	Davies	37200
141	Rajs	42000
107	Lorentz	50400
200	Whalen	52800
124	Mourgos	69600
104	Ernst	72000
202	Fay	72000
178	Grant	84000

...

20 rows selected.

Sorting by Multiple Columns

❑ The order of `ORDER BY` list is the order of sort.

```
SELECT last_name, department_id, salary
FROM    employees
ORDER BY department_id, salary DESC;
```

LAST_NAME	DEPARTMENT_ID	SALARY
Whalen	10	4400
Hartstein	20	13000
Fay	20	6000
Mourgos	50	5800
Rajs	50	3500
Davies	50	3100
Matos	50	2600
Vargas	50	2500

...

20 rows selected.

❑ You can sort by a column that is not in the `SELECT` list.
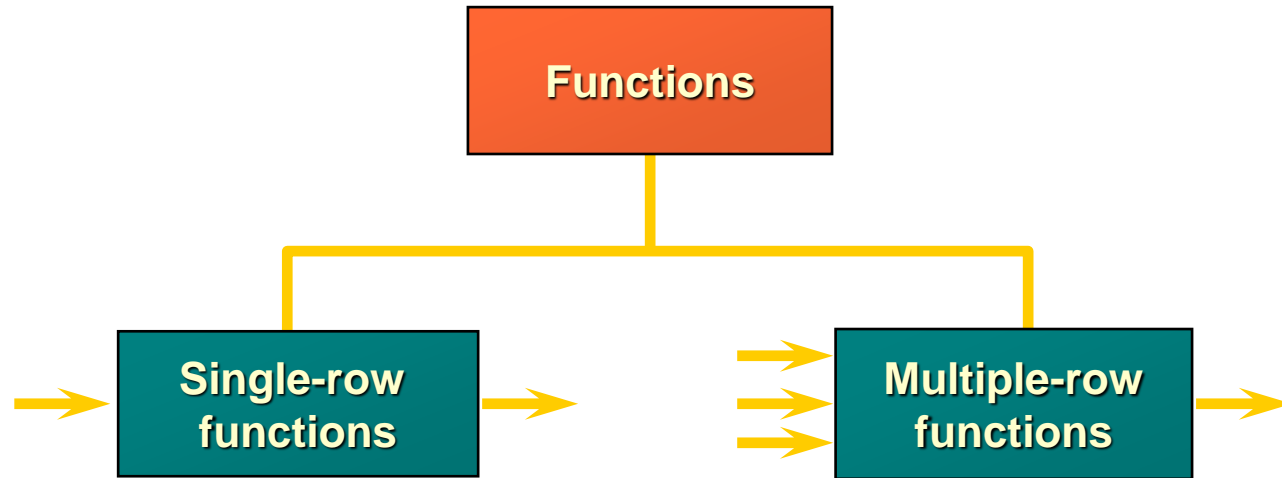
Summary

In this lesson, you should have learned how to:

❑ Use the `WHERE` clause to restrict rows of output
- ○ Use the comparison conditions
- ○ Use the `BETWEEN`, `IN`, `LIKE`, and `NULL` conditions
- ○ Apply the logical `AND`, `OR`, and `NOT` operators

❑ Use the `ORDER BY` clause to sort rows of output

```
SELECT      *|{[DISTINCT] column|expression [alias],...}
FROM        table
[WHERE      condition(s)]
[ORDER BY   {column, expr, alias} [ASC|DESC]];
```
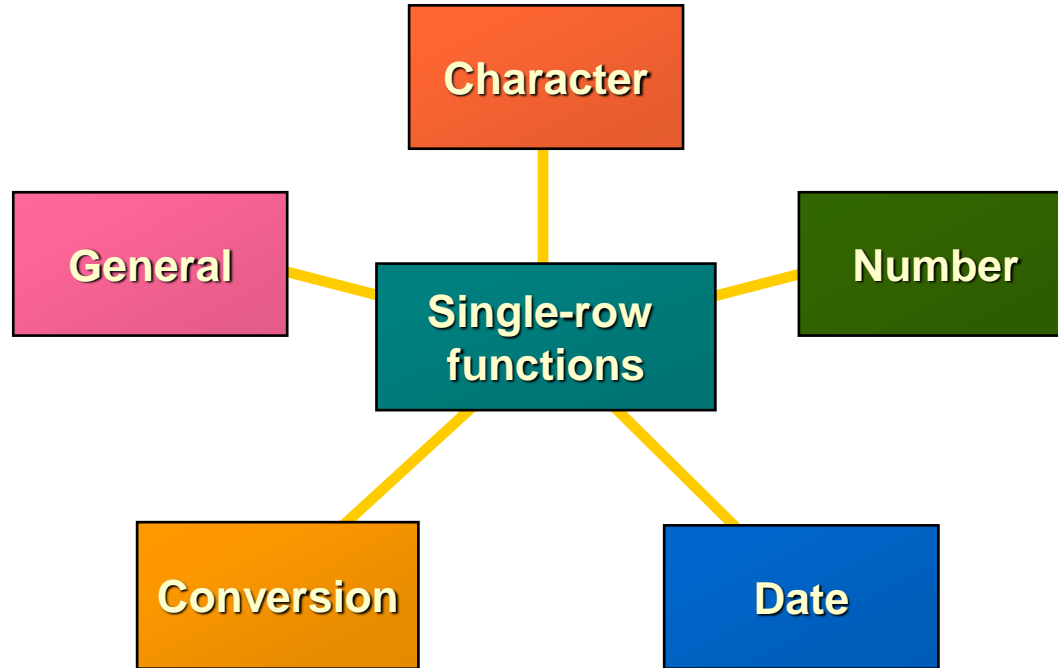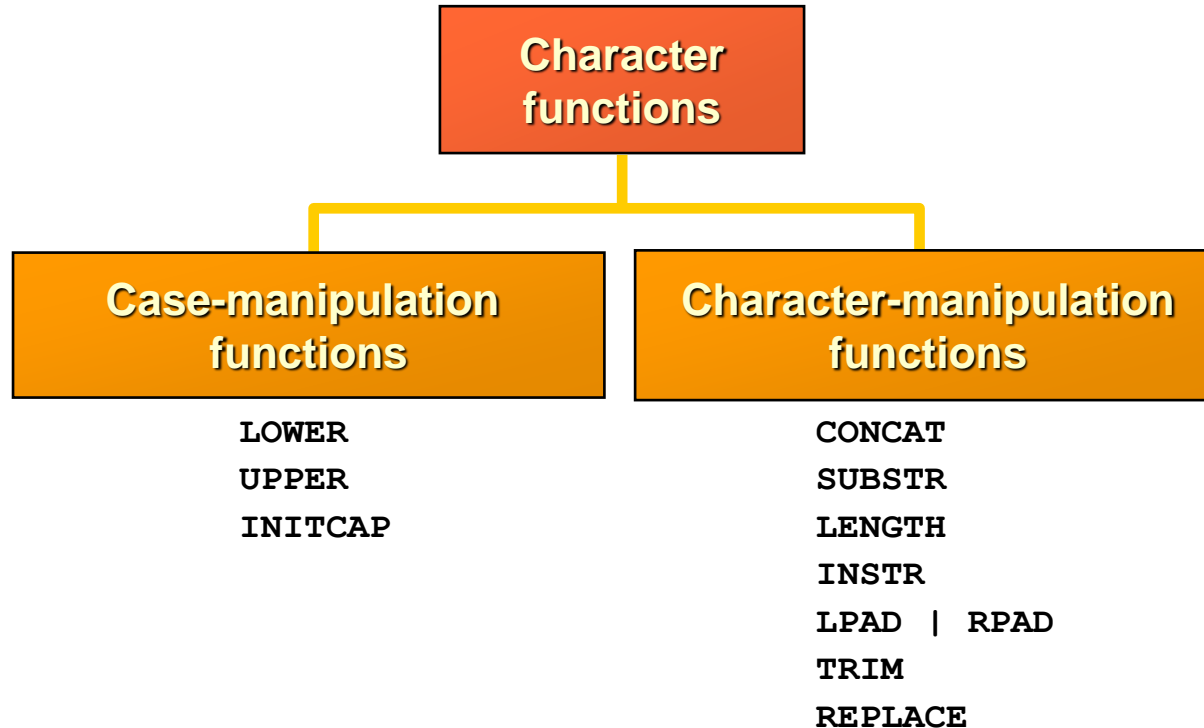
Single-Row Functions

Input

arg 1

arg 2

arg *n*

Function

Function performs action

Output

Result value

```
                    ┌─────────────────┐
                    │   Functions     │
                    └────────┬────────┘
                             │
              ┌──────────────┴──────────────┐
   ┌──────────────────┐          ┌──────────────────┐
→  │   Single-row     │ →      → │  Multiple-row    │ →
   │   functions      │        → │   functions      │
   └──────────────────┘        → └──────────────────┘
```
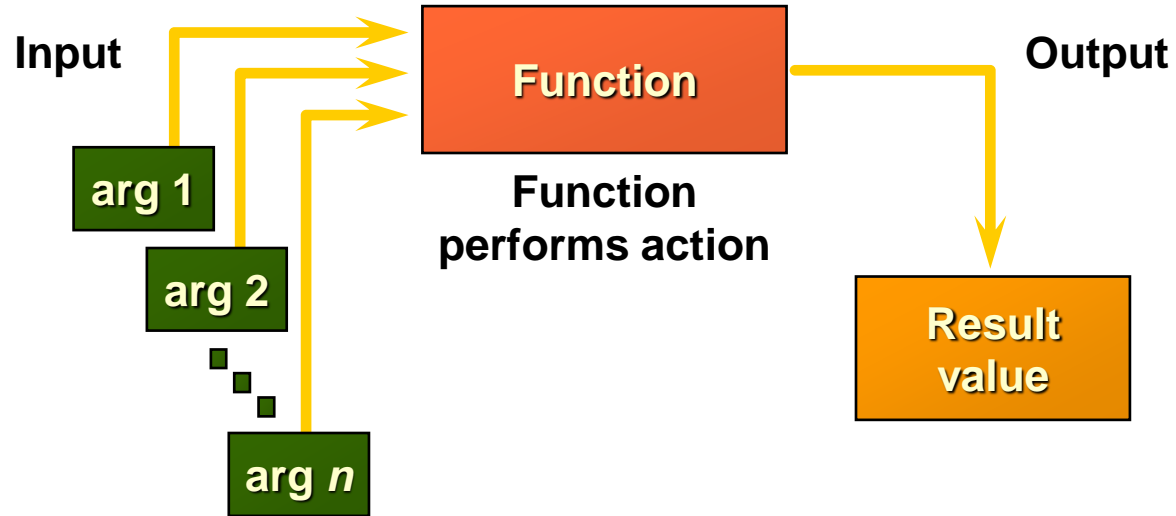
Single-Row Functions

Single row functions:

❑ Manipulate data items

❑ Accept arguments and return one value

❑ Act on each row returned

❑ Return one result per row

❑ May modify the data type

❑ Can be nested

❑ Accept arguments which can be a column or an expression

```
function_name [(arg1, arg2,...)]
```

Single-Row Functions

Character Functions

```
                    ┌─────────────────┐
                    │   Character     │
                    │   functions     │
                    └─────────────────┘
              ┌───────────────┴───────────────┐
    ┌─────────────────────┐       ┌─────────────────────────┐
    │  Case-manipulation  │       │  Character-manipulation │
    │     functions       │       │       functions         │
    └─────────────────────┘       └─────────────────────────┘
```

Case-manipulation functions

LOWER
UPPER
INITCAP

Character-manipulation functions

CONCAT
SUBSTR
LENGTH
INSTR
LPAD | RPAD
TRIM
REPLACE

Case Manipulation Functions

❑ These functions convert case for character strings.

Function	Result
`LOWER('SQL Course')`	`sql course`
`UPPER('SQL Course')`	`SQL COURSE`
`INITCAP('SQL Course')`	`Sql Course`

❑ Display the employee number, name, and department number for employee Higgins:

```
SELECT  employee_id, last_name, department_id
FROM    employees
WHERE   last_name = 'higgins';
no rows selected
```

```
SELECT  employee_id, last_name, department_id
FROM    employees
WHERE   LOWER(last_name) = 'higgins';
```

EMPLOYEE_ID	LAST_NAME	DEPARTMENT_ID
205	Higgins	110

Character-Manipulation Functions

❑ These functions manipulate character strings:

Function	Result
`CONCAT('Hello', 'World')`	`HelloWorld`
`SUBSTR('HelloWorld',1,5)`	`Hello`
`LENGTH('HelloWorld')`	`10`
`INSTR('HelloWorld', 'W')`	`6`
`LPAD(salary,10,'*')`	`*****24000`
`RPAD(salary, 10, '*')`	`24000*****`
`TRIM('H' FROM 'HelloWorld')`	`elloWorld`

```
SELECT  employee_id, CONCAT(first_name, last_name) NAME,  ①
        job_id, LENGTH (last name),  ②
        INSTR(last_name, 'a') "Contains 'a'?"  ③
FROM    employees
WHERE   SUBSTR(job_id, 4) = 'REP';
```

EMPLOYEE_ID	NAME	JOB_ID	LENGTH(LAST_NAME)	Contains 'a'?
174	EllenAbel	SA_REP	4	0
176	JonathonTaylor	SA_REP	6	2
178	KimberelyGrant	SA_REP	5	3
202	PatFay	MK_REP	3	2

Number Functions

❑ `ROUND`: Rounds value to specified decimal

   `ROUND(45.926, 2)` ⟶ `45.93`

❑ `TRUNC`:   Truncates value to specified decimal

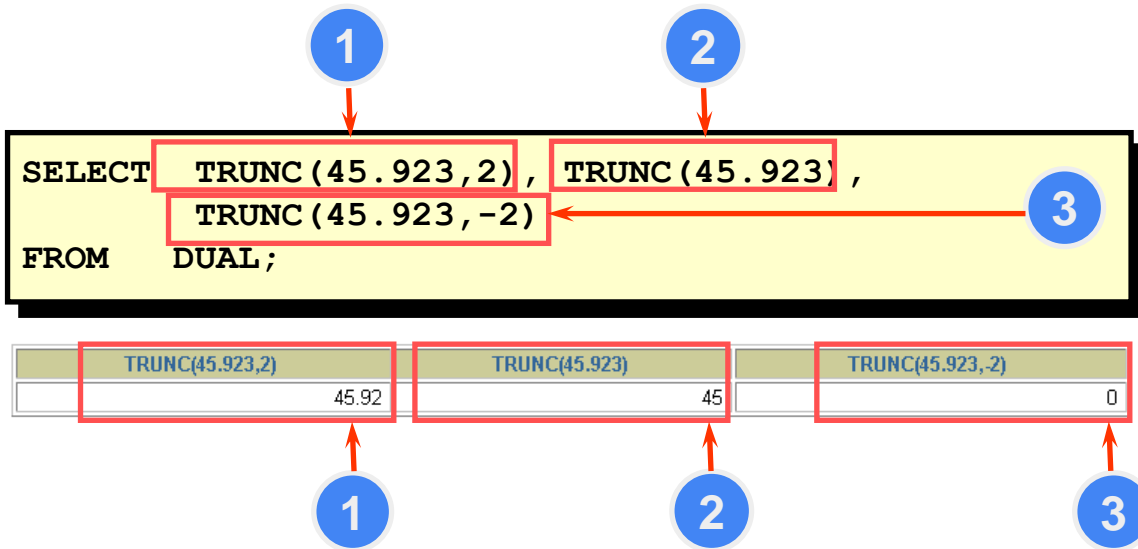   `TRUNC(45.926, 2)` ⟶ `45.92`

❑ `MOD`: Returns remainder of division

   `MOD(1600, 300)` ⟶ `100`

Using the ROUND Function

```sql
SELECT  ROUND(45.923,2),  ROUND(45.923,0),
        ROUND(45.923,-1)
FROM    DUAL;
```

ROUND(45.923,2)	ROUND(45.923,0)	ROUND(45.923,-1)
45.92	46	50

DUAL is a dummy table you can use to view results
from functions and calculations. Postgres does not need it!!

Using the TRUNC Function

```
SELECT    TRUNC(45.923,2),  TRUNC(45.923),
          TRUNC(45.923,-2)
FROM      DUAL;
```

TRUNC(45.923,2)	TRUNC(45.923)	TRUNC(45.923,-2)
45.92	45	0

❑ Calculate the remainder of a salary after it is divided by 5000 for all employees whose job title is sales representative.

```
SELECT last_name, salary, MOD(salary, 5000)
FROM    employees
WHERE   job_id = 'SA_REP';
```

LAST_NAME	SALARY	MOD(SALARY,5000)
Abel	11000	1000
Taylor	8600	3600
Grant	7000	2000

Working with Dates

Function	Description
`MONTHS_BETWEEN`	Number of months between two dates
`ADD_MONTHS`	Add calendar months to date
`NEXT_DAY`	Next day of the date specified
`LAST_DAY`	Last day of the month
`ROUND`	Round date
`TRUNC`	Truncate date
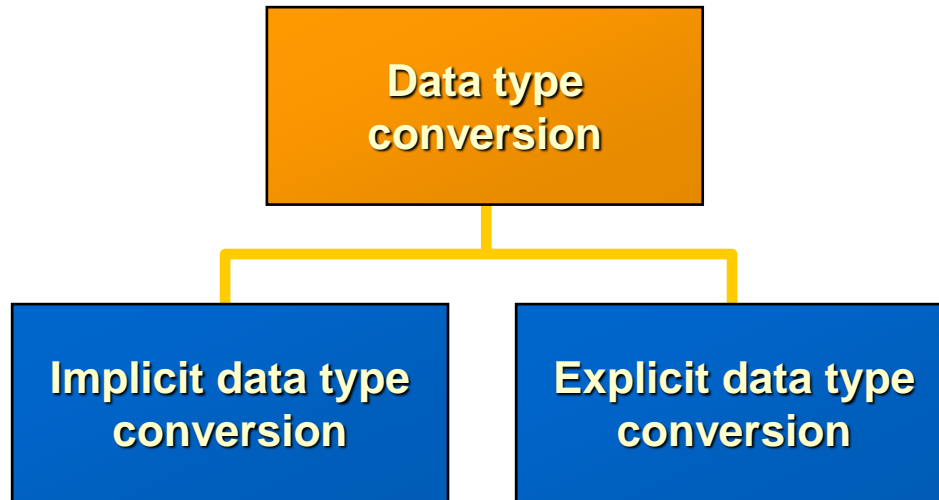
Nesting Functions

❑ Single-row functions can be nested to any level.

❑ Nested functions are evaluated from deepest level to the least deep level.

```
F3(F2(F1(col,arg1),arg2),arg3)
```

Step 1 = Result 1

Step 2 = Result 2

Step 3 = Result 3
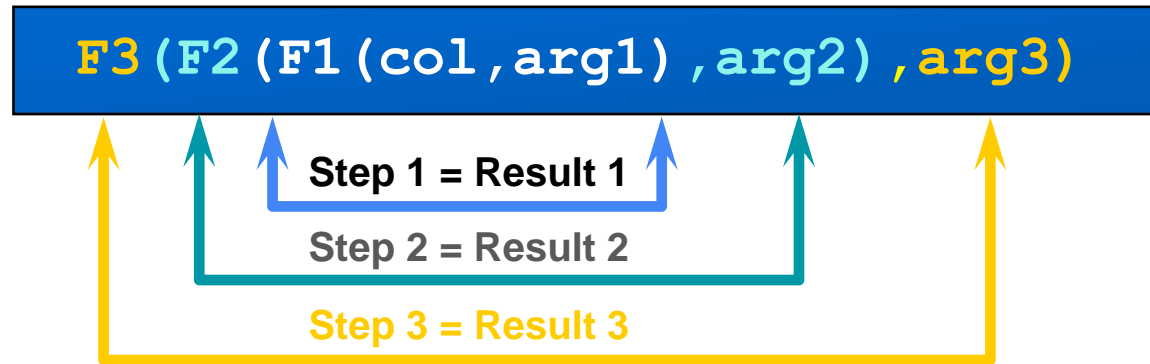
```
SELECT last_name,
        coalesce (null,null,'No Manager')
FROM    employees
WHERE   manager_id IS NULL;
```

LAST_NAME	NVL(TO_CHAR(MANAGER_ID),'NOMANAGER')
King	No Manager

Using the CASE Expression

❑ Facilitates conditional inquiries by doing the work of an IF-THEN-ELSE statement:

```
SELECT last_name, job_id, salary,
       CASE job_id WHEN 'IT_PROG'  THEN  1.10*salary
                   WHEN 'ST_CLERK' THEN  1.15*salary
                   WHEN 'SA_REP'   THEN  1.20*salary
       ELSE        salary END      "REVISED_SALARY"
FROM   employees;
```

LAST_NAME	JOB_ID	SALARY	REVISED_SALARY
...			
Lorentz	IT_PROG	4200	4620
Mourgos	ST_MAN	5800	5800
Rajs	ST_CLERK	3500	4025
...			
Gietz	AC_ACCOUNT	8300	8300

20 rows selected.

Displaying Data
from Multiple Tables

Nesting Functions

❑ Single-row functions can be nested to any level.

❑ Nested functions are evaluated from deepest level to the least deep level.

```
F3(F2(F1(col,arg1),arg2),arg3)
```

Step 1 = Result 1

Step 2 = Result 2

Step 3 = Result 3
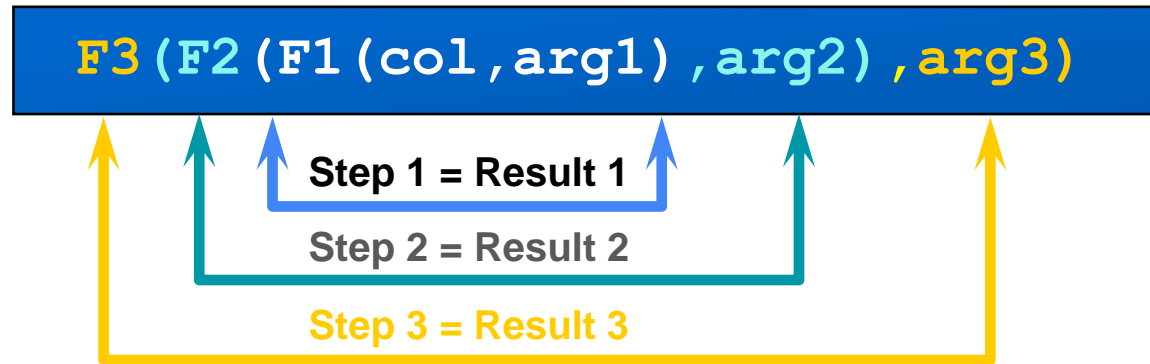
Displaying Data
from Multiple Tables

Obtaining Data from Multiple Tables

EMPLOYEES

EMPLOYEE_ID	LAST_NAME	DEPARTMENT_ID
100	King	90
101	Kochhar	90
...		
202	Fay	20
205	Higgins	110
206	Gietz	110

DEPARTMENTS

DEPARTMENT_ID	DEPARTMENT_NAME	LOCATION_ID
10	Administration	1700
20	Marketing	1800
50	Shipping	1500
60	IT	1400
80	Sales	2500
90	Executive	1700
110	Accounting	1700
190	Contracting	1700

EMPLOYEE_ID	DEPARTMENT_ID	DEPARTMENT_NAME
200	10	Administration
201	20	Marketing
202	20	Marketing
...		
102	90	Executive
205	110	Accounting
206	110	Accounting

Cartesian Products

❑ A join combines two or more tables side by side. If you do not specify how to join the tables, you get a Cartesian product. This means that SQL combines each row from the first table with every row from the second table.

❑ A Cartesian product is formed when:
  ○ A join condition is omitted
  ○ A join condition is invalid
  ○ All rows in the first table are joined to all rows in the second table

❑ To avoid a Cartesian product, always include a valid join condition in a `WHERE` clause.

Generating a Cartesian Product

❑ `SELECT A.*, B.* FROM FRUITS A, SIZES B`

Fruits
Apples
Mangoes

Sizes
Small
Medium
Big

Cartesian Product And Resultant Data

Fruits	Sizes
Apples	Small
Mangoes	Small
Apples	Medium
Mangoes	Medium
Apples	Big
Mangoes	Big

Generating a Cartesian Product

EMPLOYEES **(20 rows)**

EMPLOYEE_ID	LAST_NAME	DEPARTMENT_ID
100	King	90
101	Kochhar	90

...

EMPLOYEE_ID	LAST_NAME	DEPARTMENT_ID
202	Fay	20
205	Higgins	110
206	Gietz	110

20 rows selected.

DEPARTMENTS **(8 rows)**

DEPARTMENT_ID	DEPARTMENT_NAME	LOCATION_ID
10	Administration	1700
20	Marketing	1800
50	Shipping	1500
60	IT	1400
80	Sales	2500
90	Executive	1700
110	Accounting	1700
190	Contracting	1700

8 rows selected.

Cartesian product: 20x8=160 rows

EMPLOYEE_ID	DEPARTMENT_ID	LOCATION_ID
100	90	1700
101	90	1700
102	90	1700
103	60	1700
104	60	1700
107	60	1700

...

160 rows selected.

Join

❑ `SELECT A.fruitName, B.sizeName FROM FRUITS A,SIZES B`
`WHERE A.FRUITID = B.FRUITID;`

Results   **Messages**

	fruitName ⌄	sizeName ⌄
1	Apples	Small
2	Apples	Big
3	Mangoes	Medium

If we apply the join condition, we will get the output accordingly as given here. In this way, we can avoid Cartesian product and can get the values according to our requirements.

Joining Tables

❑ Use a join to query data from more than one table.

```
SELECT      table1.column, table2.column
FROM        table1, table2
WHERE       table1.column1 = table2.column2;
```

❑ Write the join condition in the `WHERE` clause.
❑ Prefix the column name with the table name when the same column name appears in more than one table.

What is an Equijoin?

❑ An equijoin is **a join based on equality or matching column values**. This equality is indicated with an equal sign (=) as the comparison operator in the WHERE clause, as the following query shows.

EMPLOYEES

EMPLOYEE_ID	DEPARTMENT_ID
200	10
201	20
202	20
124	50
141	50
142	50
143	50
144	50
103	60
104	60
107	60
149	80
174	80
176	80

...

DEPARTMENTS

DEPARTMENT_ID	DEPARTMENT_NAME
10	Administration
20	Marketing
20	Marketing
50	Shipping
50	Shipping
50	Shipping
50	Shipping
50	Shipping
60	IT
60	IT
60	IT
80	Sales
80	Sales
80	Sales

...

Foreign key   **Primary key**

```
SELECT  employees.employee_id, employees.last_name,
        employees.department_id, departments.department_id,
        departments.location_id
FROM    employees, departments
WHERE   employees.department_id = departments.department_id,
```

EMPLOYEE_ID	LAST_NAME	DEPARTMENT_ID	DEPARTMENT_ID	LOCATION_ID
200	Whalen	10	10	1700
201	Hartstein	20	20	1800
202	Fay	20	20	1800
124	Mourgos	50	50	1500
141	Rajs	50	50	1500
142	Davies	50	50	1500
143	Matos	50	50	1500
144	Vargas	50	50	1500

...

19 rows selected.

```
SELECT last_name, employees.department_id,department_name
 FROM    employees, departments
 WHERE   employees.department_id = departments.department_id
 AND     last_name = 'Matos'
```

EMPLOYEES

LAST_NAME	DEPARTMENT_ID
Whalen	10
Hartstein	20
Fay	20
Mourgos	50
Rajs	50
Davies	50
Matos	50
Vargas	50
Hunold	60
Ernst	60

...

DEPARTMENTS

DEPARTMENT_ID	DEPARTMENT_NAME
10	Administration
20	Marketing
20	Marketing
50	Shipping
50	Shipping
50	Shipping
50	Shipping
50	Shipping
60	IT
60	IT

...

Qualifying Ambiguous Column Names

❑ Use table prefixes to qualify column names that are in multiple tables.

❑ Improve performance by using table prefixes.

❑ Distinguish columns that have identical names but reside in different tables by using column aliases.

Using Table Aliases

❑ Simplify queries by using table aliases.

❑ Improve performance by using table prefixes.

```
SELECT  e.employee_id, e.last_name, e.department_id,
        d.department_id, d.location_id
FROM    employees e , departments d
WHERE   e.department_id = d.department_id;
```

Joining More than Two Tables

❑ To join *n* tables together, you need a minimum of n-1 join conditions. For example, to join three tables, a minimum of two joins is required.

EMPLOYEES

LAST_NAME	DEPARTMENT_ID
King	90
Kochhar	90
De Haan	90
Hunold	60
Ernst	60
Lorentz	60
Mourgos	50
Rajs	50
Davies	50
Matos	50
Vargas	50
Zlotkey	80
Abel	80
Taylor	80

...

20 rows selected.

DEPARTMENTS

DEPARTMENT_ID	LOCATION_ID
10	1700
20	1800
50	1500
60	1400
80	2500
90	1700
110	1700
190	1700

8 rows selected.

LOCATIONS

LOCATION_ID	CITY
1400	Southlake
1500	South San Francisco
1700	Seattle
1800	Toronto
2500	Oxford

Non-Equijoins

EMPLOYEES

LAST_NAME	SALARY
King	24000
Kochhar	17000
De Haan	17000
Hunold	9000
Ernst	6000
Lorentz	4200
Mourgos	5800
Rajs	3500
Davies	3100
Matos	2600
Vargas	2500
Zlotkey	10500
Abel	11000
Taylor	8600

…

20 rows selected.

JOB_GRADES

GRA	LOWEST_SAL	HIGHEST_SAL
A	1000	2999
B	3000	5999
C	6000	9999
D	10000	14999
E	15000	24999
F	25000	40000

← **Salary in the `EMPLOYEES` table must be between lowest salary and highest salary in the `JOB_GRADES` table.**

```
SELECT  e.last_name, e.salary, j.grade_level
FROM    employees e, job_grades j
WHERE   e.salary
        BETWEEN j.lowest_sal AND j.highest_sal;
```

LAST_NAME	SALARY	GRA
Matos	2600	A
Vargas	2500	A
Lorentz	4200	B
Mourgos	5800	B
Rajs	3500	B
Davies	3100	B
Whalen	4400	B
Hunold	9000	C
Ernst	6000	C

...

20 rows selected.

Self Joins

EMPLOYEES (WORKER)

EMPLOYEE_ID	LAST_NAME	MANAGER_ID
100	King	
101	Kochhar	100
102	De Haan	100
103	Hunold	102
104	Ernst	103
107	Lorentz	103
124	Mourgos	100

...

EMPLOYEES (MANAGER)

EMPLOYEE_ID	LAST_NAME
100	King
101	Kochhar
102	De Haan
103	Hunold
104	Ernst
107	Lorentz
124	Mourgos

...

MANAGER_ID in the WORKER table is equal to EMPLOYEE_ID in the MANAGER table.

Joining a Table to Itself

```
SELECT  worker.last_name || ' works for '
        || manager.last_name
FROM    employees worker, employees manager
WHERE   worker.manager_id = manager.employee_id ;
```

WORKER.LAST_NAME||'WORKSFOR'||MANAGER.LAST_NAME
Kochhar works for King
De Haan works for King
Mourgos works for King
Zlotkey works for King
Hartstein works for King
Whalen works for Kochhar
Higgins works for Kochhar
Hunold works for De Haan
Ernst works for Hunold

...
19 rows selected.

Creating Cross Joins

❑ The `CROSS JOIN` clause produces the cross-product of two tables.
❑ This is the same as a Cartesian product between the two tables.

```
SELECT last_name, department_name
FROM    employees
CROSS JOIN departments ;
```

LAST_NAME	DEPARTMENT_NAME
King	Administration
Kochhar	Administration
De Haan	Administration
Hunold	Administration

...

160 rows selected.

Creating Natural Joins

❑ The `NATURAL JOIN` clause is based on all columns in the two tables that have the same name.

❑ It selects rows from the two tables that have equal values in all matched columns.

❑ If the columns having the same names have different data types, an error is returned.

```
SELECT department_id, department_name,
       location_id, city
FROM   departments
NATURAL JOIN locations ;
```

DEPARTMENT_ID	DEPARTMENT_NAME	LOCATION_ID	CITY
60	IT	1400	Southlake
50	Shipping	1500	South San Francisco
10	Administration	1700	Seattle
90	Executive	1700	Seattle
110	Accounting	1700	Seattle
190	Contracting	1700	Seattle
20	Marketing	1800	Toronto
80	Sales	2500	Oxford

8 rows selected.

Creating Joins with the `USING` Clause

❑ If several columns have the same names but the data types do not match, the `NATURAL JOIN` clause can be modified with the `USING` clause to specify the columns that should be used for an equijoin.

❑ Use the `USING` clause to match only one column when more than one column matches.

❑ Do not use a table name or alias in the referenced columns.

❑ The `NATURAL JOIN` and `USING` clauses are mutually exclusive.

Retrieving Records with the `USING` Clause

```
SELECT e.employee_id, e.last_name, d.location_id
FROM    employees e JOIN departments d
USING (department_id) ;
```

EMPLOYEE_ID	LAST_NAME	LOCATION_ID
200	Whalen	1700
201	Hartstein	1800
202	Fay	1800
124	Mourgos	1500
141	Rajs	1500
142	Davies	1500
143	Matos	1500
144	Vargas	1500
103	Hunold	1400

...
19 rows selected.

```
SELECT employee_id, last_name,
        employees.department_id, location_id
  FROM    employees, departments
  WHERE   employees.department_id =
departments.department_id;
```

Creating Joins with the ON Clause

❑ The join condition for the natural join is basically an equijoin of all columns with the same name.

❑ To specify arbitrary conditions or specify columns to join, the ON clause is used.

❑ The join condition is separated from other *search* conditions.

❑ The ON clause makes code easy to understand.

```
SELECT  e.employee_id, e.last_name, e.department_id,
        d.department_id, d.location_id
FROM    employees e JOIN departments d
ON      (e.department_id = d.department_id);
```

EMPLOYEE_ID	LAST_NAME	DEPARTMENT_ID	DEPARTMENT_ID	LOCATION_ID
200	Whalen	10	10	1700
201	Hartstein	20	20	1800
202	Fay	20	20	1800
124	Mourgos	50	50	1500
141	Rajs	50	50	1500
142	Davies	50	50	1500
143	Matos	50	50	1500

...

19 rows selected.

```
SELECT  employee_id, city, department_name
FROM    employees e
JOIN    departments d
ON      d.department_id = e.department_id
JOIN    locations l
ON      d.location_id = l.location_id;
```

EMPLOYEE_ID	CITY	DEPARTMENT_NAME
103	Southlake	IT
104	Southlake	IT
107	Southlake	IT
124	South San Francisco	Shipping
141	South San Francisco	Shipping
142	South San Francisco	Shipping
143	South San Francisco	Shipping
144	South San Francisco	Shipping

...

19 rows selected.

DEPARTMENTS

DEPARTMENT_NAME	DEPARTMENT_ID
Administration	10
Marketing	20
Shipping	50
IT	60
Sales	80
Executive	90
Accounting	110
Contracting	190

8 rows selected.

EMPLOYEES

DEPARTMENT_ID	LAST_NAME
90	King
90	Kochhar
90	De Haan
60	Hunold
60	Ernst
60	Lorentz
50	Mourgos
50	Rajs
50	Davies
50	Matos
50	Vargas
80	Zlotkey

…

20 rows selected.

There are no employees in department 190.

Outer Joins Syntax

❑ You use an outer join to also see rows that do not meet the join condition.
❑ The left and right joint are the syntax.

```
SELECT    table1.column, table2.column
FROM      table1 left join table2
on        table1.column = table2.column;
```

```
SELECT    table1.column, table2.column
FROM      table1 right join table2
on        table1.column table2.column;
```

LEFT OUTER JOIN

```
SELECT e.last_name, e.department_id, d.department_name
FROM    employees e
LEFT OUTER JOIN departments d
ON    (e.department_id = d.department_id) ;
```

LAST_NAME	DEPARTMENT_ID	DEPARTMENT_NAME
Whalen	10	Administration
Fay	20	Marketing
Hartstein	20	Marketing

...

LAST_NAME	DEPARTMENT_ID	DEPARTMENT_NAME
De Haan	90	Executive
Kochhar	90	Executive
King	90	Executive
Gietz	110	Accounting
Higgins	110	Accounting
Grant		

20 rows selected.

RIGHT OUTER JOIN

```
SELECT e.last_name, e.department_id, d.department_name
FROM    employees e
RIGHT OUTER JOIN departments d
ON      (e.department_id = d.department_id) ;
```

LAST_NAME	DEPARTMENT_ID	DEPARTMENT_NAME
King	90	Executive
Kochhar	90	Executive

...

LAST_NAME	DEPARTMENT_ID	DEPARTMENT_NAME
Whalen	10	Administration
Hartstein	20	Marketing
Fay	20	Marketing
Higgins	110	Accounting
Gietz	110	Accounting
		Contracting

20 rows selected.

INNER Versus OUTER Joins

❑ In SQL: 1999, the join of two tables returning only matched rows is an inner join.

❑ A join between two tables that returns the results of the inner join as well as unmatched rows left (or right) tables is a left (or right) outer join.

❑ A join between two tables that returns the results of an inner join as well as the results of a left and right join is a full <u>outer</u> join.

FULL OUTER JOIN

```
SELECT e.last_name, e.department_id, d.department_name
FROM    employees e
FULL OUTER JOIN departments d
ON   (e.department_id = d.department_id) ;
```

LAST_NAME	DEPARTMENT_ID	DEPARTMENT_NAME
Whalen	10	Administration
Fay	20	Marketing

...

LAST_NAME	DEPARTMENT_ID	DEPARTMENT_NAME
De Haan	90	Executive
Kochhar	90	Executive
King	90	Executive
Gietz	110	Accounting
Higgins	110	Accounting
Grant		
		Contracting

21 rows selected.

Additional Conditions

```sql
SELECT  e.employee_id, e.last_name, e.department_id,
        d.department_id, d.location_id
FROM    employees e JOIN departments d
ON      (e.department_id = d.department_id)
AND     e.manager_id = 149 ;
```

EMPLOYEE_ID	LAST_NAME	DEPARTMENT_ID	DEPARTMENT_ID	LOCATION_ID
174	Abel	80	80	2500
176	Taylor	80	80	2500

Aggregating Data
Using Group Functions

Maryam Ramezani

What Are Group Functions?

Group functions operate on sets of rows to give one result per group.

EMPLOYEES

DEPARTMENT_ID	SALARY
90	24000
90	17000
90	17000
60	9000
60	6000
60	4200
50	5800
50	3500
50	3100
50	2600
50	2500
80	10500
80	11000
80	8600
	7000
10	4400

...

20 rows selected.

The maximum salary in the EMPLOYEES table.

MAX(SALARY)
24000

Types of Group Functions (Aggregations)

- ❑ AVG
- ❑ COUNT
- ❑ MAX
- ❑ MIN
- ❑ STDDEV
- ❑ SUM
- ❑ VARIANCE

Group Functions Syntax

```
SELECT [column,] group_function(column), ...
FROM            table
[WHERE condition]
[GROUP BY        column]
[ORDER BY        column];
```

Using the `AVG` and `SUM` Functions

❑ You can use `AVG` and `SUM` for numeric data.

```
SELECT  AVG(salary), MAX(salary),
        MIN(salary), SUM(salary)
FROM    employees
WHERE   job_id LIKE '%REP%';
```

AVG(SALARY)	MAX(SALARY)	MIN(SALARY)	SUM(SALARY)
8150	11000	6000	32600

Using the `MIN` and `MAX` Functions

❑ You can use `MIN` and `MAX` for any data type.

```
SELECT MIN(hire_date), MAX(hire_date)
FROM    employees;
```

MIN(HIRE_	MAX(HIRE_
17-JUN-87	29-JAN-00

```
SELECT MIN(last name), MAX(last name)
   FROM    employees;
```

Using the `COUNT` Function

❑ `COUNT(*)` returns the number of rows in a table.

```
SELECT  COUNT(*)
FROM    employees
WHERE   department_id = 50;
```

COUNT(*)
5

Using the COUNT Function

❑ COUNT(*expr*) returns the number of rows with non-null values for the *expr*.

❑ Display the number of department values in the EMPLOYEES table, excluding the null values.

```
SELECT COUNT(commission_pct)
FROM   employees
WHERE  department_id = 80;
```

COUNT(COMMISSION_PCT)
3

Using the `DISTINCT` Keyword

❑ `COUNT(DISTINCT expr)` returns the number of distinct non-null values of the *expr*.

❑ Display the number of distinct department values in the `EMPLOYEES` table.

```
SELECT  COUNT(DISTINCT department_id)
FROM    employees;
```

COUNT(DISTINCTDEPARTMENT_ID)
7

Group Functions and Null Values

❑ Group functions ignore null values in the column.

```
select avg(opening_amt) from customer
```

❑ The coalesce function forces group functions to include null values.

```
select avg(coalesce(opening_amt,0)) from customer
```

EMPLOYEES

DEPARTMENT_ID	SALARY
10	4400
20	13000
20	6000
50	5800
50	3500
50	3100
50	2500
50	2600
60	9000
60	6000
60	4200
80	10500
80	8600
80	11000
90	24000
90	17000

...

20 rows selected.

4400

9500

3500

6400

10033

The average salary in EMPLOYEES table for each department.

DEPARTMENT_ID	AVG(SALARY)
10	4400
20	9500
50	3500
60	6400
80	10033.3333
90	19333.3333
110	10150
	7000

❑ Divide rows in a table into smaller groups by using the `GROUP BY` clause.

```
SELECT column, group_function(column)
FROM            table
[WHERE condition]
[GROUP BY        group_by_expression]
[ORDER BY        column];
```

❑ All columns in the `SELECT` list that are not in group functions must be in the `GROUP BY` clause.

```
SELECT     department_id, AVG(salary)
FROM       employees
GROUP BY department_id ;
```

DEPARTMENT_ID	AVG(SALARY)
10	4400
20	9500
50	3500
60	6400
80	10033.3333
90	19333.3333
110	10150
	7000

8 rows selected.

❑ The `GROUP BY` column does not have to be in the `SELECT` list.

```
SELECT    AVG(salary)
FROM      employees
GROUP BY department_id ;
```

AVG(SALARY)
4400
9500
3500
6400
10033.3333
19333.3333
10150
7000

Grouping by More Than One Column

EMPLOYEES

DEPARTMENT_ID	JOB_ID	SALARY
90	AD_PRES	24000
90	AD_VP	17000
90	AD_VP	17000
60	IT_PROG	9000
60	IT_PROG	6000
60	IT_PROG	4200
50	ST_MAN	5800
50	ST_CLERK	3500
50	ST_CLERK	3100
50	ST_CLERK	2600
50	ST_CLERK	2500
80	SA_MAN	10500
80	SA_REP	11000
80	SA_REP	8600

...

DEPARTMENT_ID	JOB_ID	SALARY
20	MK_REP	6000
110	AC_MGR	12000
110	AC_ACCOUNT	8300

20 rows selected.

"Add up the salaries in the EMPLOYEES table for each job, grouped by department.

DEPARTMENT_ID	JOB_ID	SUM(SALARY)
10	AD_ASST	4400
20	MK_MAN	13000
20	MK_REP	6000
50	ST_CLERK	11700
50	ST_MAN	5800
60	IT_PROG	19200
80	SA_MAN	10500
80	SA_REP	19600
90	AD_PRES	24000
90	AD_VP	34000
110	AC_ACCOUNT	8300
110	AC_MGR	12000
	SA_REP	7000

13 rows selected.

```
SELECT    department_id dept_id, job_id, SUM(salary)
FROM      employees
GROUP BY department_id, job_id ;
```

DEPT_ID	JOB_ID	SUM(SALARY)
10	AD_ASST	4400
20	MK_MAN	13000
20	MK_REP	6000
50	ST_CLERK	11700
50	ST_MAN	5800
60	IT_PROG	19200
80	SA_MAN	10500
80	SA_REP	19600
90	AD_PRES	24000
90	AD_VP	34000
110	AC_ACCOUNT	8300
110	AC_MGR	12000
	SA_REP	7000

13 rows selected.

Illegal Queries – Using Group Functions

❑ Any column or expression in the `SELECT` list that is not an aggregate function must be in the `GROUP BY` clause.

❑ Column missing in the `GROUP BY` clause:

```
SELECT  department_id, COUNT(last_name)
FROM    employees;
```

```
SELECT  department_id, COUNT(last_name)
        *
ERROR at line 1:
ORA-00937: not a single-group group function
```

Illegal Queries – Using Group Functions

❏ You cannot use the `WHERE` clause to restrict groups.

❏ You use the `HAVING` clause to restrict groups.

❏ You cannot use group functions in the `WHERE` clause.

❏ Cannot use the WHERE clause to restrict groups:

```
SELECT    department_id, AVG(salary)
FROM      employees
WHERE     AVG(salary) > 8000
GROUP BY department_id;
```

```
WHERE   AVG(salary) > 8000
        *
ERROR at line 3:
ORA-00934: group function is not allowed here
```

EMPLOYEES

DEPARTMENT_ID	SALARY
90	24000
90	17000
90	17000
60	9000
60	6000
60	4200
50	5800
50	3500
50	3100
50	2600
50	2500
80	10500
80	11000
80	8600

...

20	6000
110	12000
110	8300

20 rows selected.

The maximum salary per department when it is greater than $10,000

DEPARTMENT_ID	MAX(SALARY)
20	13000
80	11000
90	24000
110	12000

Use the `HAVING` clause to restrict groups:

1. Rows are grouped.
2. The group function is applied.
3. Groups matching the `HAVING` clause are displayed.

```
SELECT column, group_function
FROM             table
[WHERE condition]
[GROUP BY        group_by_expression]
[HAVING          group_condition]
[ORDER BY        column];
```

```
SELECT     department_id, MAX(salary)
FROM       employees
GROUP BY department_id
HAVING     MAX(salary)>10000 ;
```

DEPARTMENT_ID	MAX(SALARY)
20	13000
80	11000
90	24000
110	12000

Using the `HAVING` Clause

```
SELECT    job_id, SUM(salary) PAYROLL
FROM      employees
WHERE     job_id NOT LIKE '%REP%'
GROUP BY job_id
HAVING    SUM(salary) > 13000
ORDER BY SUM(salary);
```

JOB_ID	PAYROLL
IT_PROG	19200
AD_PRES	24000
AD_VP	34000

Nesting Group Functions

Display the maximum average salary.

```
SELECT    MAX(AVG(salary))
FROM      employees
GROUP BY department_id;
```

MAX(AVG(SALARY))
19333.3333

Summary

You should have learned how to:

- ❑ Use the group functions `COUNT, MAX, MIN, AVG`
- ❑ Write queries that use the `GROUP BY` clause
- ❑ Write queries that use the `HAVING` clause

```
SELECT column, group_function(column)
FROM               table
[WHERE condition]
[GROUP BY          group_by_expression]
[HAVING            group_condition]
[ORDER BY          column];
```

Subqueries

Who has a salary greater than Abel's?

Main Query:

Which employees have salaries greater than Abel's salary?

Subquery

What is Abel's salary?

Nested Queries – Definitions

❑ A nested query is a query inside another query
  ○ The enclosing query also called outer query
  ○ Nested query is called inner query
❑ It usually appears as a condition in where or having clauses.
❑ There can be multiple levels of nesting
❑ There are two kinds of nested queries
  ○ Correlated
  ○ Non-Correlated

Example:
```
Select movie_title
From movies
Where director_id IN (
            Select person_id
            From People
            Where person_state = 'TX')
```

Nested Queries: Non-Correlated

❑ Generates data required by outer query before it can be executed
❑ Inner query does not contain any reference to outer query
❑ Behaves like a procedure
❑ The result should not contain any column from the nested query
❑ Example

Schema:
o People(person_fname, person_lname, person_id, person_state, person_city)
o Movies(movie_id, movie_title, director_id, studio_id)

Query: Select movie_title, studio_id
           From Movies
           Where director_id IN (Select person_id
                       From People
                       Where person_state = 'TX')

Steps:
o Subquery is executed
o Subquery results are plugged into the outer query
o The outer query is processed

Nested Queries: Correlated

- ❑ Contains reference to the outer query
- ❑ Behaves like a loop
  ## Example:

Schema: `People(person_fname, person_lname, person_id, person_state, person_city)`
`Cast_Movies(cast_member_id, role, movie_id)`

Query: `select person_fname, person_lname`
`From People p1`
`Where 'Pam Green' in ( Select role`
`From Cast_Movies`
`Where p1.person_id = cast_member_id)`

Steps:
- ○ Contents of the table row in outer query are read
- ○ Sub-query is executed using data in the row being processed.
- ○ Results of the inner query are passed to the where in the outer query
- ○ The Outer query is Processed
- ○ Loop continues till all rows are exhausted

Subquery Syntax

❑ The subquery (inner query) executes once before the main query.
❑ The result of the subquery is used by the main query (outer query).

```
SELECT      select_list
FROM        table
WHERE       expr operator
                  (SELECT          select_list
                         FROM table);
```

Using a Subquery

```
SELECT  last_name
FROM    employees        11000
WHERE   salary >
                 (SELECT  salary
                  FROM     employees
                  WHERE    last_name = 'Abel');
```

LAST_NAME
King
Kochhar
De Haan
Hartstein
Higgins

Guidelines for Using Subqueries

❑ Enclose subqueries in parentheses.

❑ Place subqueries on the right side of the comparison condition.

❑ The `ORDER BY` clause in the subquery is not needed unless you are performing Top-N analysis.

❑ Use single-row operators with single-row subqueries and use multiple-row operators with

multiple-row subqueries.

Types of Subqueries

□ Types of Subqueries
- **Single-row subqueries**: Queries that return only one row from the inner `SELECT` statement
- **Multiple-row subqueries**: Queries that return more than one row from the inner `SELECT` statement
  - **Note:** There are also **multiple-column subqueries**: Queries that return more than one column from the inner `SELECT` statement.

- **Single-row subquery**

Main query	
	Subquery

returns ⟶ **ST_CLERK**

- **Multiple-row subquery**

Main query	
	Subquery

returns ⟶ **ST_CLERK**
SA_MAN

Single-Row Subqueries
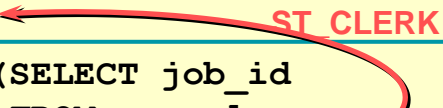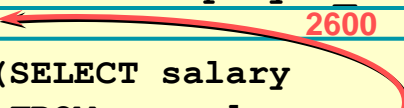
❑ Return only one row

❑ Use single-row comparison operators

Operator	Meaning
=	Equal to
>	Greater than
>=	Greater than or equal to
<	Less than
<=	Less than or equal to
<>	Not equal to

Executing Single-Row Subqueries
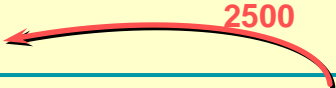
```
SELECT last_name, job_id, salary
FROM    employees
WHERE   job_id =                ST_CLERK
                (SELECT job_id
                 FROM     employees
                 WHERE    employee_id = 141)
AND     salary >                2600
                (SELECT salary
                 FROM     employees
                 WHERE    employee_id = 143);
```

LAST_NAME	JOB_ID	SALARY
Rajs	ST_CLERK	3500
Davies	ST_CLERK	3100

Using Group Functions in a Subquery

```
SELECT last_name, job_id, salary
FROM    employees              2500
WHERE   salary =
                (SELECT MIN(salary)
                 FROM    employees);
```

LAST_NAME	JOB_ID	SALARY
Vargas	ST_CLERK	2500

❑ The Oracle server executes subqueries first.

❑ The Oracle server returns results into the `HAVING` clause of the main query.

```
SELECT    department_id, MIN(salary)
FROM      employees
GROUP BY department_id
HAVING    MIN(salary) >
                            (SELECT MIN(salary)
                             FROM    employees
                             WHERE   department_id = 50);
```

2500

What is Wrong with this Statement?

```
SELECT  employee_id, last_name
FROM    employees
WHERE   salary =
                (SELECT    MIN(salary)
                 FROM      employees
                 GROUP BY department_id);
```

```
ERROR at line 4:
ORA-01427: single-row subquery returns more than
one row
```

Single-row operator with multiple-row subquery

Will this Statement Return Rows?

```
SELECT last_name, job_id
FROM    employees
WHERE   job_id =
                (SELECT job_id
                 FROM    employees
                 WHERE   last_name = 'Haas');
```

```
no rows selected
```

Subquery returns no values

Multiple-Row Subqueries

❑ Return more than one row

❑ Use multiple-row comparison operators

Operator	Meaning
IN	Equal to any member in the list
ANY	Compare value to each value returned by the subquery
ALL	Compare value to every value returned by the subquery

```
SELECT  employee_id, last_name, job_id, salary
FROM    employees          9000, 6000, 4200
WHERE   salary < ANY
                    (SELECT  salary
                     FROM    employees
                     WHERE   job_id = 'IT_PROG')

AND     job_id <> 'IT_PROG';
```

EMPLOYEE_ID	LAST_NAME	JOB_ID	SALARY
124	Mourgos	ST_MAN	5800
141	Rajs	ST_CLERK	3500
142	Davies	ST_CLERK	3100
143	Matos	ST_CLERK	2600
144	Vargas	ST_CLERK	2500

...
10 rows selected.

```
SELECT  employee_id, last_name, job_id, salary
FROM    employees
WHERE   salary < ALL              9000, 6000, 4200

                  (SELECT salary
                   FROM    employees
                   WHERE   job_id = 'IT_PROG')
AND     job_id <> 'IT_PROG';
```

EMPLOYEE_ID	LAST_NAME	JOB_ID	SALARY
141	Rajs	ST_CLERK	3500
142	Davies	ST_CLERK	3100
143	Matos	ST_CLERK	2600
144	Vargas	ST_CLERK	2500

```
SELECT  emp.last_name
FROM    employees emp
WHERE   emp.employee_id NOT IN
                            (SELECT mgr.manager_id
                             FROM    employees mgr);


no rows selected
```

$$\left\{ \begin{array}{ll} \textbf{theta} & \textbf{ANY} \\ \textbf{theta} & \textbf{SOME} \\ \textbf{theta} & \textbf{ALL} \end{array} \right\} \qquad \text{theta} \in \left\{ \begin{array}{c} = \\ \neq \\ < \\ \leq \\ \geq \\ > \end{array} \right\}$$

- The ALL keyword modifies the greater than comparison operator to mean greater than all values.
- The ANY keyword is not as restrictive as the ALL keyword.
- When used with the greater than comparison operator, "> ANY" means greater than some value.
- The "= ANY" operator is exactly equivalent to the IN operator.
- However, the "!= ANY" (not equal any) is not equivalent to the NOT IN operator.

Any, Some, All

❑ Give the providers whose status are not maximum.

1- SELECT   S#

   FROM   S

   WHERE   STATUS < **ANY (SELECT DISTINCT**   STATUS   **FROM**   S)

2- SELECT   S#

   FROM   S

   WHERE   STATUS < **(SELECT   MAX** (STATUS)   **FROM**   S)

Subqueries and the EXISTS Operator

❑ When a subquery uses the EXISTS operator, the subquery functions as an existence test.

❑ The WHERE clause of the outer query tests for the existence of rows returned by the inner query.

❑ The subquery does not actually produce any data; rather, it returns a value of TRUE or FALSE.

❑ The general format of a subquery WHERE clause with an EXISTS operator is shown here.

❑ Note that the NOT operator can also be used to negate the result of the EXISTS operator.

```
WHERE [NOT] EXISTS (subquery)
```

Example

```
SELECT emp_last_name "Last Name", emp_first_name "First Name"
FROM employee
WHERE EXISTS
     (SELECT *
      FROM dependent
      WHERE emp_ssn = dep_emp_ssn);
```

```
Last Name  First Name
---------- ----------------
Joyner     Suzanne
Zhu        Waiman
Bock       Douglas
```

Subqueries and the EXISTS operator

❑ Subqueries using an EXISTS operator are a bit different from other subqueries, in the following ways:

- o The keyword EXISTS is not preceded by a column name, constant, or other expression.
- o The SELECT clause list of a subquery that uses an EXISTS operator almost always consists of an asterisk (∗). This is because there is no real point in listing column names since you are simply testing for the existence of rows that meet the conditions specified in the subquery.
- o The subquery evaluates to TRUE or FALSE rather than returning any data.
- o A subquery that uses an EXISTS operator will always be a correlated subquery.

Subqueries and the EXISTS operator

❑ The EXISTS operator is very important, because there is often no alternative to its use.

❑ All queries that use the IN operator or a modified comparison operator (=, <, >, etc. modified by ANY or ALL) can be expressed with the EXISTS operator.

❑ However, some queries formulated with EXISTS cannot be expressed in any other way!

❑ The NOT EXISTS operator is the mirror-image of the EXISTS operator.

❑ A query that uses NOT EXISTS in the WHERE clause is satisfied if the subquery returns no rows.

Subqueries and the EXISTS operator

```
SELECT emp_last_name
FROM employee
WHERE emp_ssn = ANY
    (SELECT dep_emp_ssn
     FROM dependent);
```

```
EMP_LAST_NAME
-------------
Bock
Zhu
Joyner
```

```
SELECT
    emp_last_name
FROM employee
WHERE EXISTS
   (SELECT *
    FROM dependent
    WHERE emp_ssn
   = dep_emp_ssn);
```

```
EMP_LAST_NAME
-----------------
Bock
Zhu
Joyner
```

Subqueries and the ORDER BY Clause

❑ The SELECT statement shown below adds the ORDER BY clause to specify sorting by first name within last name.

❑ Note that the ORDER BY clause is placed after the WHERE clause, and that this includes the subquery as part of the WHERE clause.

```
SELECT emp_last_name "Last Name",
 emp_first_name "First Name"
FROM employee
WHERE EXISTS
    (SELECT *
     FROM dependent
     WHERE emp_ssn = dep_emp_ssn)
ORDER BY emp_last_name, emp_first_name;
```

Output:

```
Last Name   First Name
----------  ----------
Bock        Douglas
Joyner      Suzanne
Zhu         Waiman
```

Union

❑ Union Joins allow multiple query results to be combined into a single result set

Syntax
```
Select select_list
From table [,table, ….]
[Where condition]
Union [All]
Select select_list
From table [,table, ….]
[Where condition]
```

Example
```
Select person_id,
person_city, person_state
  From People
  Union
  Select studio_id,
      studio_city,
      studio_state
  From Studios
```

❑ Notes:
  ○ The number of columns selected for both the queries should be the same
  ○ The columns are merged in order in which they are selected
  ○ The duplicates are eliminated from the combined table
  ○ More than two tables can be joined together

Union (All & Order By)

❑ **Union query eliminates all duplicates in the resultant table**
  ○ Union All is used when we do not want to eliminate the duplicates
  ○ Union and Union distinct are the same.
❑ **Union and Order By can be used together to order the results of the combined table**
  ○ This clause is not allowed when a single column result is obtained and the all keyword is used since the duplicates are eliminated and there is nothing to order by
❑ **Example**
```
Select studio_id, studio_state
From Studios
Union
Select Person_id, person_state
From People
Order By studio_state
```

Intersect

❑ In the Intersect Query results of two separate queries are concatenated, however, only common elements of the two queries are included in the resultset

❑ Example

```
Select person_state
From People
Intersect
Select studio_state
From Studios
```

SELECT R.A
FROM R, S
WHERE R.A=S.A
EXCEPT
SELECT R.A
FROM R, T
WHERE R.A=T.A

Q_1 Q_2

Subquery Benefits

❑ They can simplify the logic and readability of your query, especially if you need to filter or aggregate data before joining it with another table.

❑ They can help you avoid duplicate rows or columns that might result from a join operation.

❑ They can enable you to perform complex calculations or comparisons that might not be possible with a join.

  ○ For example, you can use a subquery to find the average salary of each department, and then compare it with the salary of each employee in the main query.

Subquery Drawbacks

Subqueries also have some drawbacks that can affect database performance.

❑ They can increase the processing time and memory usage of your query, especially if the subquery returns a large number of rows or columns.

❑ They can limit the optimization options of the database system, as some subqueries cannot use indexes or other techniques to speed up the execution.

❑ They can introduce errors or inconsistencies if the subquery is not correlated with the main query, or if the subquery data changes during the execution of the main query.

Join Benefits

Joins are another way to query data from multiple tables in a database.

❏ They can reduce the number of queries and subqueries needed to retrieve the data you want, which can save processing time and memory.

❏ They can leverage the indexes and other features of the database system to optimize the join operation and make it faster and more efficient.

❏ They can ensure the consistency and accuracy of the data, as the join condition determines which rows from each table are matched and returned.

Join Drawbacks

❏ They can complicate the syntax and readability of your query, especially if you need to join multiple tables or use different types of joins.

❏ They can generate unwanted or redundant rows or columns that might affect the quality and size of the result set.

❏ They can require careful planning and design of the database schema and the join condition, as poorly structured or indexed tables or columns can slow down or fail the join operation.

How to choose

❑ Deciding whether to use a subquery or a join for your query is dependent on various factors, such as the data structure, the query complexity, the database system, and the performance goals.

❑ As a general guideline, you should use a subquery if you need to filter or aggregate data before joining it with another table, or if you need to perform calculations or comparisons that are not possible with a join. On the other hand, if you need to query data from multiple tables based on a common column or condition, or if you want to take advantage of the optimization features of the database system, then using a join is recommended.

❑ Ultimately, it is best to test and compare the execution time and result set of both options and choose the one that meets your requirements and expectations.

Nested queries as alternatives to INTERSECT and EXCEPT

```
(SELECT R.A, R.B
 FROM   R)
INTERSECT
(SELECT S.A, S.B
 FROM   S)
```

```
SELECT R.A, R.B
FROM    R
WHERE  EXISTS(
            SELECT *
            FROM S
            WHERE R.A=S.A
AND R.B=S.B)
```

INTERSECT and EXCEPT
not in some DBMSs!

```
(SELECT R.A, R.B
 FROM   R)
EXCEPT
(SELECT S.A, S.B
 FROM   S)
```

```
SELECT R.A, R.B
FROM    R
WHERE  NOT EXISTS(
            SELECT *
            FROM S
            WHERE R.A=S.A AND
R.B=S.B)
```

If R, S have no duplicates, then
can write without sub-queries
(HOW?)

Manipulating Data

Adding a New Row to a Table

70	Public Relations	100	1700	**New row**

DEPARTMENTS

DEPARTMENT_ID	DEPARTMENT_NAME	MANAGER_ID	LOCATION_ID
10	Administration	200	1700
20	Marketing	201	1800
50	Shipping	124	1500
60	IT	103	1400
80	Sales	149	2500
90	Executive	100	1700
110	Accounting	205	1700
190	Contracting		1700

…insert a new row into the DEPARMENTS table…

DEPARTMENT_ID	DEPARTMENT_NAME	MANAGER_ID	LOCATION_ID
10	Administration	200	1700
20	Marketing	201	1800
50	Shipping	124	1500
60	IT	103	1400
80	Sales	149	2500
90	Executive	100	1700
110	Accounting	205	1700
190	Contracting		1700
70	Public Relations	100	1700

The `INSERT` Statement Syntax

❑ Add new rows to a table by using the `INSERT` statement.

❑ Only one row is inserted at a time with this syntax.

```
INSERT INTO table [(column [, column...])]
VALUES              (value [, value...]);
```

Inserting New Rows

❑ Insert a new row containing values for each column.

❑ List values in the default order of the columns in the table.

❑ Optionally, list the columns in the `INSERT` clause.

```
INSERT INTO departments(department_id, department_name,
                        manager_id, location_id)
VALUES        (70, 'Public Relations', 100, 1700);
1 row created.
```

❑ Enclose character and date values within single quotation marks.

Inserting Rows with Null Values

❑ Implicit method: Omit the column from the column list.

```
INSERT INTO departments (department_id,
                              department_name  )
VALUES          (30, 'Purchasing');
1 row created.
```

- **Explicit method: Specify the `NULL` keyword in the `VALUES` clause.**

```
INSERT INTO departments
VALUES          (100, 'Finance', NULL, NULL);
1 row created.
```

Inserting Special Values

❑ The `current_date` function records the current date.

```
INSERT INTO employees (employee_id,
                first_name, last_name,
                email, phone_number,
                hire_date, job_id, salary,
                commission_pct, manager_id,
                department_id)
VALUES          (113,
                'Louis', 'Popp',
                'LPOPP', '515.124.4567',
                current_date, 'AC_ACCOUNT', 6900,
                NULL, 205, 100);
1 row created.
```

❑ Use : substitution in a SQL statement to prompt for values.

❑ : is a placeholder for the variable value.

```
INSERT INTO departments
            (department_id, department_name, location_id)
VALUES      (:department_id, ':department_name', :location);
```

Define Substitution Variables

"department_id" | 40
"department_name" | Human Resources
"location" | 2500

Submit for Execution   Cancel

```
1 row created.
```

Copying Rows from Another Table

❑ Write your `INSERT` statement with a subquery.

```
INSERT INTO sales_reps(id, name, salary, commission_pct)
  SELECT employee_id, last_name, salary, commission_pct
  FROM    employees
  WHERE   job_id LIKE '%REP%';

4 rows created.
```

❑ Do not use the `VALUES` clause.

❑ Match the number of columns in the `INSERT` clause to those in the subquery.

Changing Data in a Table

EMPLOYEES

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	HIRE_DATE	JOB_ID	SALARY	DEPARTMENT_ID	COMMISSION_P
100	Steven	King	SKING	17-JUN-87	AD_PRES	24000	90	
101	Neena	Kochhar	NKOCHHAR	21-SEP-89	AD_VP	17000	90	
102	Lex	De Haan	LDEHAAN	13-JAN-93	AD_VP	17000	90	
103	Alexander	Hunold	AHUNOLD	03-JAN-90	IT_PROG	9000	60	
104	Bruce	Ernst	BERNST	21-MAY-91	IT_PROG	6000	60	
107	Diana	Lorentz	DLORENTZ	07-FEB-99	IT_PROG	4200	60	
124	Kevin	Mourgos	KMOURGOS	16-NOV-99	ST_MAN	5800	50	

Update rows in the EMPLOYEES table.

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	HIRE_DATE	JOB_ID	SALARY	DEPARTMENT_ID	COMMISSIO
100	Steven	King	SKING	17-JUN-87	AD_PRES	24000	90	
101	Neena	Kochhar	NKOCHHAR	21-SEP-89	AD_VP	17000	90	
102	Lex	De Haan	LDEHAAN	13-JAN-93	AD_VP	17000	90	
103	Alexander	Hunold	AHUNOLD	03-JAN-90	IT_PROG	9000	30	
104	Bruce	Ernst	BERNST	21-MAY-91	IT_PROG	6000	30	
107	Diana	Lorentz	DLORENTZ	07-FEB-99	IT_PROG	4200	30	
124	Kevin	Mourgos	KMOURGOS	16-NOV-99	ST_MAN	5800	50	

❑ Modify existing rows with the `UPDATE` statement.

```
UPDATE              table
SET                 column = value [, column = value, ...]
[WHERE              condition];
```

❑ Update more than one row at a time, if required.

Updating Rows in a Table

❑ Specific row or rows are modified if you specify the `WHERE` clause.

```
UPDATE  employees
SET     department id = 70
WHERE   employee_id = 113;
1 row updated.
```

❑ All rows in the table are modified if you omit the `WHERE` clause.

```
UPDATE      copy_emp
SET         department_id = 110;
22 rows updated.
```

Updating Two Columns with a Subquery

❑ Update employee 114's job and salary to match that of employee 205.

```
UPDATE    employees
SET       job_id  = (SELECT  job_id
                     FROM     employees
                     WHERE    employee_id = 205),
          salary  = (SELECT  salary
                     FROM     employees
                     WHERE    employee_id = 205)
WHERE     employee_id   =  114;
1 row updated.
```

Updating Rows Based on Another Table

❑ Use subqueries in `UPDATE` statements to update rows in a table based on values from another table.

```
UPDATE   copy_emp
SET      department_id  =  (SELECT department_id
                            FROM employees
                            WHERE employee_id = 100)
WHERE    job_id         =  (SELECT job_id
                            FROM employees
                            WHERE employee_id = 200);
1 row updated.
```

Updating Rows: Integrity Constraint Error

Department number 55 does not exist in the parent table!

```
UPDATE  employees
SET     department_id = 55
WHERE   department_id = 110;
```

```
UPDATE  employees
        *
ERROR at line 1:
ORA-02291: integrity constraint (HR.EMP_DEPT_FK)
violated - parent key not found
```

DEPARTMENTS

DEPARTMENT_ID	DEPARTMENT_NAME	MANAGER_ID	LOCATION_ID
10	Administration	200	1700
20	Marketing	201	1800
30	Purchasing		
100	Finance		
50	Shipping	124	1500
60	IT	103	1400

Delete a row from the DEPARTMENTS table.

DEPARTMENT_ID	DEPARTMENT_NAME	MANAGER_ID	LOCATION_ID
10	Administration	200	1700
20	Marketing	201	1800
30	Purchasing		
50	Shipping	124	1500
60	IT	103	1400

The `DELETE` Statement

❑ You can remove existing rows from a table by using the `DELETE` statement.

```
DELETE  [FROM]      table
[WHERE   condition];
```

Deleting Rows from a Table

❑ Specific rows are deleted if you specify the `WHERE` clause.

```
 DELETE FROM departments
 WHERE   department_name = 'Finance';
1 row deleted.
```

❑ All rows in the table are deleted if you omit the `WHERE` clause.

```
DELETE FROM  copy_emp;
22 rows deleted.
```

Deleting Rows Based on Another Table

❑ Use subqueries in `DELETE` statements to remove rows from a table based on values from another table.

```
DELETE FROM employees
WHERE   department_id =
                        (SELECT department_id
                         FROM    departments
                         WHERE   department_name LIKE '%Public%');
1 row deleted.
```

You cannot delete a row that contains a primary key that is used as a foreign key in another table.

```
DELETE FROM departments
WHERE        department_id = 60;
```

```
DELETE FROM departments
            *
ERROR at line 1:
ORA-02292: integrity constraint (HR.EMP_DEPT_FK)
violated - child record found
```

Overview of the Explicit Default Feature

❑ With the explicit default feature, you can use the `DEFAULT` keyword as a column value where the column default is desired.

❑ The addition of this feature is for compliance with the SQL: 1999 Standard.

❑ This allows the user to control where and when the default value should be applied to data.

❑ Explicit defaults can be used in `INSERT` and `UPDATE` statements.

Using Explicit Default Values

❑ `DEFAULT` with `INSERT`:

```
INSERT INTO departments
   (department_id, department_name, manager_id)
VALUES (300, 'Engineering', DEFAULT);
```

❑ `DEFAULT` with `UPDATE`:

```
UPDATE departments
SET manager_id = DEFAULT WHERE department_id = 10;
```

❑ If no default value for the corresponding column has been specified, Postgres sets the column to null.

The `MERGE` Statement

❑ Provides the ability to conditionally update or insert data into a database table

❑ **Performs an `UPDATE` if the row exists, and an `INSERT` if it is a new row:**
  ○ Avoids separate updates
  ○ Increases performance and ease of use
  ○ Is useful in data warehousing applications: you may need to work with data coming from multiple sources, some of which may be duplicates. With the `MERGE` statement, you can conditionally add or modify rows.

❑ The `MERGE` statement is deterministic. You cannot update the same row of the target table multiple times in the same `MERGE` statement.

The `MERGE` Statement Syntax

❑ You can conditionally insert or update rows in a table by using the `MERGE` statement.

```
MERGE INTO table_name table_alias
  USING (table|view|sub_query) alias
  ON (join condition)
  WHEN MATCHED THEN
    UPDATE SET
    col1 = col_val1,
    col2 = col2_val
  WHEN NOT MATCHED THEN
    INSERT (column_list)
    VALUES (column_values);
```

Merging Rows

❑ Insert or update rows in the `COPY_EMP` table to match the `EMPLOYEES` table.

    ○ The example shown matches the `EMPLOYEE_ID` in the `COPY_EMP` table to the `EMPLOYEE_ID` in the `EMPLOYEES` table. If a match is found, the row in the `COPY_EMP` table is updated to match the row in the `EMPLOYEES` table. If the row is not found, it is inserted into the `COPY_EMP` table.

```
MERGE INTO copy_emp  c
   USING employees e
   ON (c.employee_id = e.employee_id)
WHEN MATCHED THEN
   UPDATE SET
      c.first_name      = e.first_name,
      c.last_name       = e.last_name,
      ...
      c.department_id   = e.department_id
WHEN NOT MATCHED THEN
   INSERT VALUES(e.employee_id, e.first_name, e.last_name,
               e.email, e.phone_number, e.hire_date, e.job_id,
               e.salary, e.commission_pct, e.manager_id,
               e.department_id);
```

- ❑ The condition c.employee_id = e.employee_id is evaluated. Because the COPY_EMP table is empty, the condition returns false: there are no matches. The logic falls into the WHEN NOT MATCHED clause, and the MERGE command inserts the rows of the EMPLOYEES table into the COPY_EMP table.
- ❑ If rows existed in the COPY_EMP table and employee IDs matched in both tables (the COPY_EMP and EMPLOYEES tables), the existing rows in the COPY_EMP table would be updated to match the EMPLOYEES table.

```
SELECT *
FROM COPY_EMP;

no rows selected
```

```
MERGE INTO copy_emp c
  USING employees e
  ON (c.employee_id = e.employee_id)
WHEN MATCHED THEN
  UPDATE SET
     ...
WHEN NOT MATCHED THEN
  INSERT VALUES...;
```

```
SELECT *
FROM COPY_EMP;

20 rows selected.
```

Statement	Description
INSERT	Adds a new row to the table
UPDATE	Modifies existing rows in the table
DELETE	Removes existing rows from the table
MERGE	Conditionally inserts or updates data in a table