



Database Programming with SQL

9-2

Using Rollup and Cube Operations, and Grouping Sets



Objectives

This lesson covers the following objectives:

- Use ROLLUP to produce subtotal values
- Use CUBE to produce cross-tabulation values
- Use GROUPING SETS to produce a single result set
- Use the GROUPING function to identify the extra row values created by either a ROLLUP or CUBE operation

Purpose

- Let's develop the problem that you were presented with in the last lesson just a little further.
- To find the average height of all students, you use this query:

```
SELECT AVG(height) FROM students;
```

- If you want to know the average height of the students based on their years in school, you could write a number of different SQL statements like this:

```
SELECT AVG(height) FROM students WHERE year_in_school = 10;
```

```
SELECT AVG(height) FROM students WHERE year_in_school = 11;
```

```
SELECT AVG(height) FROM students WHERE year_in_school = 12;
```

Purpose

- Or you could simplify the problem by writing just one statement containing the GROUP BY and HAVING clauses.
- What if, once you have selected your groups and computed your aggregates across these groups, you also wanted subtotals per group and a grand total of all the rows selected.

Purpose

- You could import the results into a spreadsheet application, get out your calculator, or compute the totals manually on paper using arithmetic.
- But better still, you could use some of the extensions to the GROUP BY clause specifically created for this purpose: ROLLUP, CUBE, and GROUPING SETS.
- Using these extensions requires less work on your part and they are all highly efficient to use, from the point of view of the database.

ROLLUP

- In GROUP BY queries, you are quite often required to produce subtotals and totals, and the ROLLUP operation can do that for you.
- Without using the ROLLUP operator, that kind of requirement would mean writing several queries and then entering the results in, for instance, a spreadsheet to calculate and format the results.
- ROLLUP creates subtotals that roll up from the most detailed level to a grand total, using the grouping list specified in the GROUP BY clause.

ROLLUP

- The action of ROLLUP is straightforward: it creates subtotals that roll up from the most detailed level to a grand total
- ROLLUP uses an ordered list of grouping columns in its argument list.
- First, it calculates the standard aggregate values specified in the GROUP BY clause.
- Next, it creates progressively higher-level subtotals, moving from right to left through the list of grouping columns.
- Finally, it creates a grand total.

ROLLUP Result Table

- In the result table below, the rows highlighted in red are generated by the ROLLUP operation:

```
SELECT department_id, job_id, SUM(salary)
FROM employees
WHERE department_id < 50
GROUP BY ROLLUP (department_id, job_id);
```

DEPARTMENT_ID	JOB_ID	SUM(SALARY)
10	AD_ASST	4400
10	-	4400
20	MK_MAN	13000
20	MK_REP	6000
20	-	19000
-	-	23400

← Subtotal for dept_id 10

← Subtotal for dept_id 20

← Grand Total for report

ROLLUP Result Formula

- The number of columns or expressions that appear in the ROLLUP argument list determines the number of groupings.
- The formula is $(\text{number of columns}) + 1$ where number of columns is the number of columns listed in the ROLLUP argument list.



ROLLUP Result Formula

- In the example query below, two columns are listed in the ROLLUP argument list and, therefore, you will see three values generated automatically.

```
SELECT department_id, job_id, SUM(salary)
FROM employees
WHERE department_id < 50
GROUP BY ROLLUP (department_id, job_id);
```

Without ROLLUP

- If you use GROUP BY without ROLLUP for the same query, what would the results look like?

```
SELECT department_id, job_id, SUM(salary)
FROM   employees
WHERE  department_id < 50
GROUP BY (department_id, job_id);
```

DEPARTMENT_ID	JOB_ID	SUM(SALARY)
20	MK_MAN	13000
10	AD_ASST	4400
20	MK_REP	6000

- You would have to execute multiple queries to get the subtotals produced by ROLLUP.

CUBE

- CUBE, like ROLLUP, is an extension to the GROUP BY clause.
- It produces cross-tabulation reports.
- It can be applied to all aggregate functions including AVG, SUM, MIN, MAX, and COUNT.
- Columns listed in the GROUP BY clause are cross-referenced to create a superset of groups.
- The aggregate functions specified in the SELECT list are applied to this group to create summary values for the additional super-aggregate rows.

CUBE

- Every possible combination of rows is aggregated by CUBE.
- If you have n columns in the GROUP BY clause, there will be 2^n possible super-aggregate combinations.
- Mathematically these combinations form an n -dimensional cube, which is how the operator got its name.



CUBE

- CUBE is often used in queries that use columns from separate tables rather than separate columns from a single table.
- Imagine, for example, a user querying the Sales table for a company like AMAZON.COM.
- A commonly requested cross-tabulation report might include subtotals for all possible combinations of sales across a Month, Region, and Product.

CUBE

- In the following statement, the rows in red are generated by the CUBE operation:

```
SELECT department_id, job_id, SUM(salary)
FROM employees
WHERE department_id < 50
GROUP BY CUBE (department_id, job_id);
```

Total for report →

Subtotal for MK_MAN →

Subtotal for MK_REP →

Subtotal for AD_ASST →

Subtotal for dept 10 →

Subtotal for dept 20 →

DEPARTMENT_ID	JOB_ID	SUM(SALARY)
-	-	23400
-	MK_MAN	13000
-	MK_REP	6000
-	AD_ASST	4400
10	-	4400
10	AD_ASST	4400
20	-	19000
20	MK_MAN	13000
20	MK_REP	6000

GROUPING SETS

- GROUPING SETS is another extension to the GROUP BY clause.
- It is used to specify multiple groupings of data.
- It gives you the functionality of having multiple GROUP BY clauses in the same SELECT statement, which is not allowed in the normal syntax.



GROUPING SETS

- If you want to see data from the EMPLOYEES table grouped by (department_id, job_id, manager_id).
- But also grouped by (department_id, manager_id).
- And also grouped by (job_id, manager_id), then you would normally have to write three different select statements with the only difference between them being the GROUP BY clauses.



GROUPING SETS

- For the database, this means retrieving the same data three different times, and that can be quite a big overhead.
- Imagine if your company had 3,000,000 employees.
- Then you are asking the database to retrieve 9 million rows instead of just 3 million rows – quite a big difference.
- So GROUPING SETS are much more efficient when writing complex reports.

GROUPING SETS

- In the following statement, the rows highlighted in color are generated by the GROUPING SETS operation:

```
SELECT department_id, job_id, manager_id, SUM(salary)
FROM employees
WHERE department_id < 50
GROUP BY GROUPING SETS
((job_id, manager_id), (department_id, job_id), (department_id, manager_id));
```

DEPARTMENT_ID	JOB_ID	MANAGER_ID	SUM(SALARY)
-	MK_MAN	100	13000
-	MK_MAN	201	6000
-	AD_ASST	101	4400
10	AD_ASST	-	4400
20	MK_MAN	-	13000
20	MK_REP	-	6000
10	-	101	19000
20	-	100	13000
20	-	201	6000

GROUPING Functions

- When you use ROLLUP or CUBE to create reports with subtotals, you quite often also have to be able to tell which rows in the output are actual rows returned from the database and which rows are computed subtotal rows resulting from the ROLLUP or CUBE operations.

DEPARTMENT_ID	JOB_ID	SUM(SALARY)
-	-	23400
-	MK_MAN	13000
-	MK_REP	6000
-	AD_ASST	4400
10	-	4400
10	AD_ASST	4400
20	-	19000
20	MK_MAN	13000
20	MK_REP	6000

GROUPING Functions

- If you look at the report on the right, how will you be able to differentiate between the actual database rows and the calculated rows?
- How can you tell the difference between a stored NULL value returned by the query and NULL values created by a ROLLUP or CUBE.

DEPARTMENT_ID	JOB_ID	SUM(SALARY)
-	-	23400
-	MK_MAN	13000
-	MK_REP	6000
-	AD_ASST	4400
10	-	4400
10	AD_ASST	4400
20	-	19000
20	MK_MAN	13000
20	MK_REP	6000

GROUPING Functions

- The GROUPING function handles these problems.
- Using a single column from the query as its argument, the GROUPING function will return a 1 for an aggregated (computed) row and a 0 for a non-aggregated (returned) row.
- The syntax for the GROUPING is simply GROUPING (column_name).
- It is used only in the SELECT clause and it takes only a single column expression as the argument.

GROUPING Functions

- Example:

```
SELECT department_id, job_id, SUM(salary),  
       GROUPING(department_id) AS "Dept sub total",  
       GROUPING(job_id) AS "Job sub total"  
FROM employees  
WHERE department_id < 50  
GROUP BY CUBE (department_id, job_id);
```

DEPARTMENT_ID	JOB_ID	SUM(SALARY)	Dept sub total	Job sub total
-	-	23400	1	1
-	MK_MAN	13000	1	0
-	MK_REP	6000	1	0
-	AD_ASST	4400	1	0
10	-	4400	0	1
10	AD_ASST	4400	0	0
20	-	19000	0	1
20	MK_MAN	13000	0	0
20	MK_REP	6000	0	0

Terminology

Key terms used in this lesson included:

- CUBE
- GROUPING FUNCTION
- GROUPING SETS
- ROLLUP

Summary

In this lesson, you should have learned how to:

- Use ROLLUP to produce subtotal values
- Use CUBE to produce cross-tabulation values
- Use GROUPING SETS to produce a single result set
- Use the GROUPING function to identify the extra row values created by either a ROLLUP or CUBE operation

