

# Joins

## **Agenda**

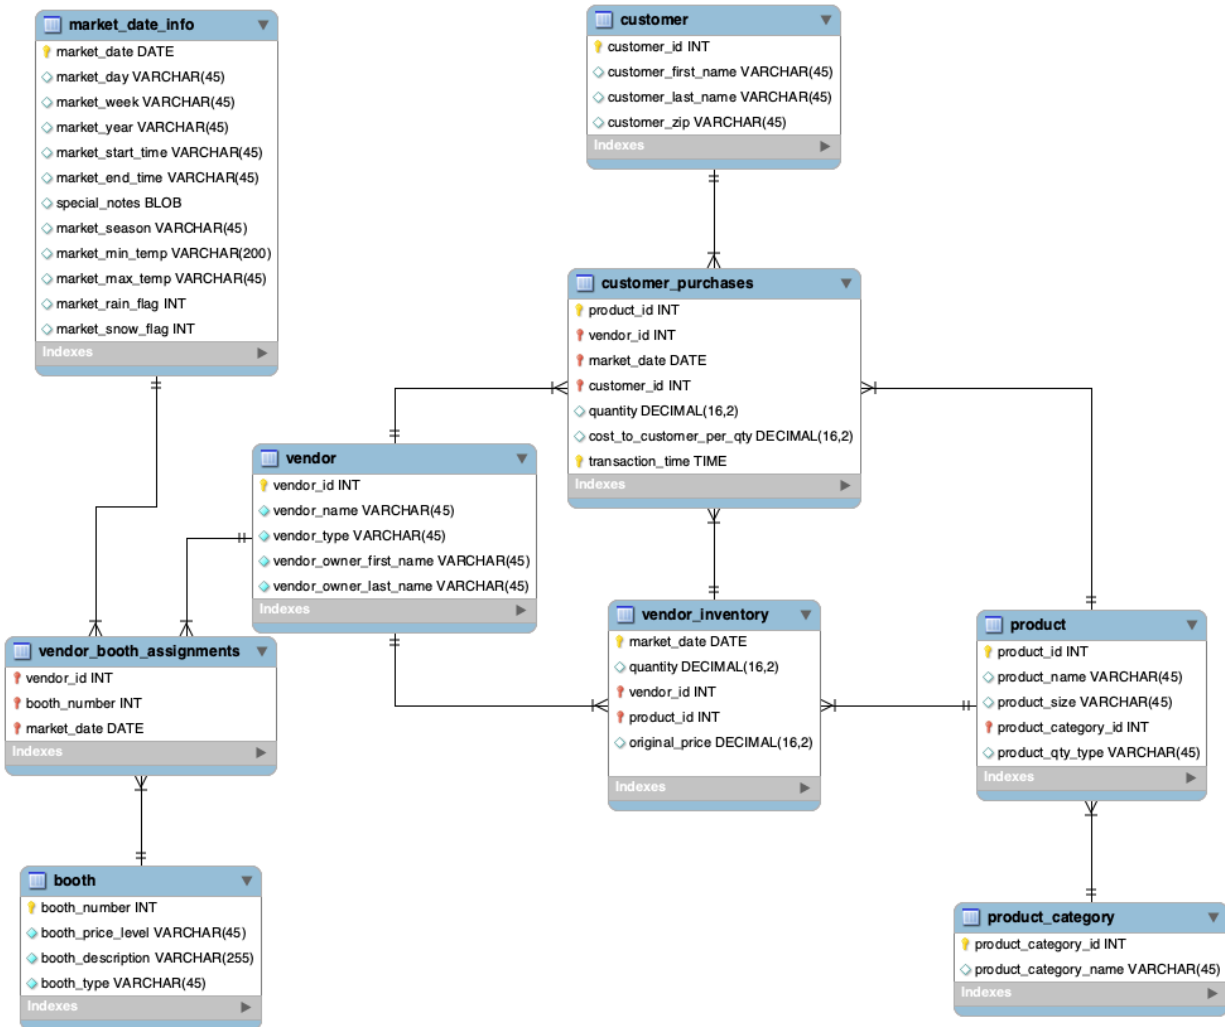
### Problem Statement

1. Joins
    - a. INNER
    - b. LEFT
    - c. RIGHT
    - d. FULL OUTER
- 

### Problem Statement:

You are a Data Analyst at Amazon Fresh. You have been tasked to study the Farmer's Market.

Dataset: Farmer's Market database



So far...

- You have learned to select data from **a single database table** and filter to the rows you want.
- But you might wonder what to do if the data you need exists across multiple related tables in the database.

Question: Get details of all vendors selling products along with the name of each product they sell and the quantity of that product present in their inventory?

## Intuition:

Now, what tables do we require?

1. to get the vendor details like vendor name & type - the **vendor** table.
2. to get the details about each specific item, including product name & size - the **product** table.
3. to find out the quantity of the product for each vendor - the **vendor\_inventory** table.

So you need all this combined information from 3 tables here.

This is where **SQL JOINS** come in.

## Solution Query:


```
SELECT
    v.vendor_name, v.vendor_type
    p.product_name, p.product_size
    vi.quantity
FROM farmers_market.vendor v
JOIN farmers_market.vendor_inventory vi
    ON v.vendor_id = vi.vendor_id
JOIN farmers_market.product p
    ON vi.product_id = p.product_id
ORDER BY v.vendor_name;
```

Note that ER Diagrams are crucial in identifying which tables we can join and which key fields connect them.

Question: List all the products along with their product category name.

Since only the ID of the product category exists in the **product** table, and the product category's name is in the **product\_category** table,

We have to combine the data in the product and product\_category tables together to generate this list.




- The figure shows the one-to-many relationship between these tables.
- Their primary keys are each identified with an asterisk and the foreign key with a double asterisk.
- Each row in the **product\_category** table can be associated with many rows in the product table, but each row in the **product** table is associated with only one row in the **product\_category** table.
- The fields that connect the two tables are `product_category.product_category_id` and `product.product_category_id`.

**Now, there are multiple ways of joining these tables.**

*Instructor Note: [You can pick your own example or watch the lecture to understand how to explain Joins]*

# Joins in MySQL



## Syntax for joining tables:

```
SELECT [columns to return]
FROM [left table]
[JOIN TYPE] [right table]
ON [left table].[field in left table to match] = [right table].[field in
right table to match]
```

## Order of Execution of a SQL query:

- **FROM** - The database gets the data from tables in FROM clause and if necessary, performs the JOINS.
- **JOIN** - Depending on the type of JOIN used in the query and conditions specified for joining the tables in the **ON** clause, the database engine matches rows from the virtual table created in the FROM clause.
- **WHERE** - After the JOIN operation, the data is filtered based on the conditions specified in the WHERE clause. Rows that do not meet the criteria are excluded.

- **GROUP BY** - If the query includes a GROUP BY clause, the rows are grouped based on the specified columns and aggregate functions are applied to the groups created.
- **HAVING** - The HAVING clause filters the groups of rows based on the specified conditions
- **SELECT** - After grouping and filtering is done, the SELECT statement determines which columns to include in the final result set.
- **ORDER BY** - It allows you to sort the result set based on one or more columns, either in ascending or descending order.
- **OFFSET** - The specified number of rows are skipped from the beginning of the result set.
- **LIMIT** - After skipping the rows, the LIMIT clause is applied to restrict the number of rows returned.


Question: Get a list of customers' zip codes for customers who made a purchase on 2019-04-06.

Will you need a join here?


What type of join are we going to use here?

If we need zip codes of all the customers who made a purchase, we only require an **intersection of customers** whose details are present in the customer tables.

### Inner JOIN



An INNER JOIN only returns records that have matches in both tables.



### Query:

```
SELECT
    DISTINCT cp.customer_id,
    c.customer_zip
FROM farmers_market.customer c
INNER JOIN farmers_market.customer_purchases cp
ON cp.customer_id = c.customer_id
WHERE
    cp.market_date='2019-04-06';
```

### Breaking down **Inner JOIN** :

- An inner join is a type of join in SQL that returns only the rows from both tables that have matching values in the specified columns.
- In this case, the inner join is performed on the "customer" and "customer\_purchases" tables using the "customer\_id" column as the matching column.
- The "INNER JOIN" clause specifies that we want to retrieve only the rows that have matching values in both tables. In other words,

we only want to retrieve the details of customers who have made a purchase on the specified date.

- The "ON" clause specifies the condition for the join. In this case, we want to match the "customer\_id" column in both tables. By joining on this column, we can link each purchase to the corresponding customer.
- Once the join is performed, we can retrieve the zip codes of the customers who made a purchase on the specified date by selecting the "customer\_zip" column from the "customer" table.
- The "DISTINCT" keyword is used to ensure that we only get one row per customer, even if they made multiple purchases on the specified date.

---

### Quiz - 1

Q. Both of the queries will give the same result.

- `Select t1.*, t2.* from DB.tbl1 t1 join DB.tbl2 t2 on t1.col1 = t2.col1`
- `Select t1.*, t2.* from DB.tbl2 t2 join DB.tbl1 t1 on t2.col1 = t1.col1`

a) True-correct

b) False


Keyword: Both of the queries

Select Instructor: Sai Nischitha Thatha

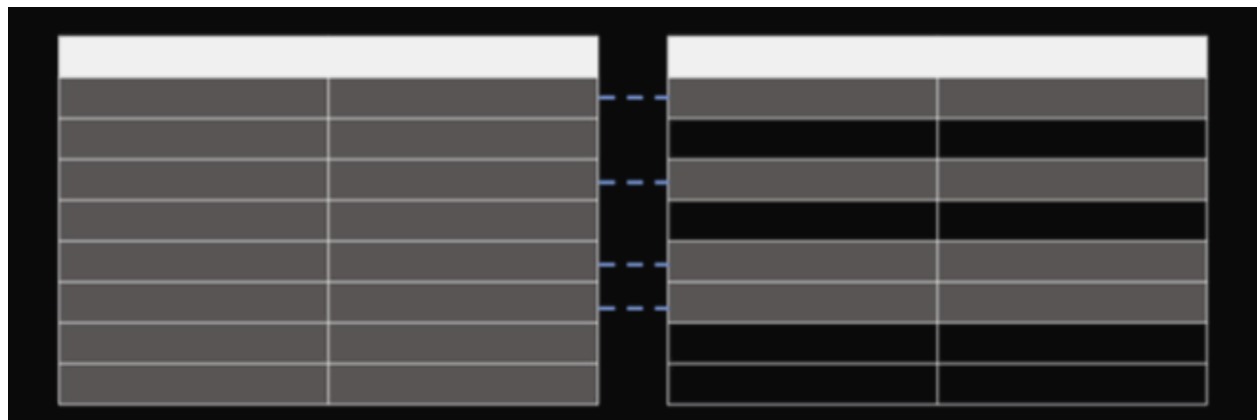
---

### Left JOIN





This tells the DBMS to pull all records from the table on the “left side” of the JOIN, and only the matching records (based on the criteria specified in the JOIN clause) from the table on the “right side” of the JOIN.



Question: As per our question we want to list ALL the products and their product categories.

**Ques.** Which table should we use as the left table if we use LEFT JOIN?

**Ans:** The product table should be on the left and product\_categories should be on the right.

**Actual Query:**

```
SELECT * FROM  
farmers_market.product  
LEFT JOIN farmers_market.product_category
```

```
ON product.product_category_id =  
product_category.product_category_id
```

**NOTE:** You may have noticed two columns called **product\_category\_id** in the output.

**That is because we selected all fields using the asterisk(\*), and there are fields in both tables with the same name.**

To remedy this, **we could either specify the list of fields to be returned** and only include the **product\_category\_id** from one of the tables or **alias the column names to indicate which table each came from.**

**Query:**

```
SELECT  
    p.product_id,  
    p.product_name,  
    pc.product_category_id,  
    pc.product_category_name  
FROM farmers_market.product AS p  
LEFT JOIN farmers_market.product_category AS pc  
ON p.product_category_id = pc.product_category_id  
ORDER BY pc.product_category_name, p.product_name;
```


Breaking down **Left JOIN** :

- The Left JOIN indicates that we want all rows from the **product** table (which is listed on the left side of the JOIN keyword) and
- only the associated rows from the **product\_category** table. So, if a category is not associated with any products, it will not be


included in the results.

- If a product were without a category, it would be included in the results, with the fields on the **product\_category** side being NULL.
- The ON part of the JOIN clause tells the query to match up the rows in the two tables using each table's values in the product\_category\_id field.
- We can specify which table each column is from since it's possible to have identically named columns in different tables.

## **Right JOIN**



In a RIGHT JOIN, all of the rows from the “right table” are returned, along with only the matching rows from the “left table,” using the fields specified in the ON part of the query.



To write the same query (the one we saw earlier) using a RIGHT JOIN, you can simply reverse the order of the tables and use a RIGHT JOIN instead of a LEFT JOIN.

### Query:

```
SELECT
  p.product_id,
  p.product_name,
  pc.product_category_id,
  pc.product_category_name
FROM farmers_market.product_category AS pc
RIGHT JOIN farmers_market.product AS p
ON p.product_category_id = pc.product_category_id
ORDER BY pc.product_category_name, p.product_name;
```

- In this query, the "product\_category" table is on the left side of the RIGHT JOIN, and the "product" table is on the right side.
- The rest of the query remains the same as in your original LEFT JOIN query.
- This RIGHT JOIN query will return all product categories, including those without associated products, and it will display products when they have a matching category.

Question: Find out the customers who are either new to the market or have deleted their account from the market.

What information is required to answer this?

1. For customers who are new to the market, we'd need the **`customer`** table because that's where the new customers are (those who haven't made any purchase yet).
2. For the customers that have left, they can only be found in the purchase history i.e. the **`customer\_purchases`** table as their records are deleted from the **`customer`** table.

---

### Quiz - 2

Q.Any result that is achieved with right join can also be achieved with left join by just swapping the tables

- a) True - correct
- b) False

Keyword: with the right join

Select Instructor: Sai Nischitha Thatha

---

### Quiz - 3

Q. Which join is used to return the book list which only Jim has read (Jim: Table A and Vinny: table B)

- a) inner
- b) Left - correct
- c) Right
- d) outer

Keyword: Which join is used

Select Instructor: Sai Nischitha Thatha

---

Q1: Get all the customers who haven't purchased anything from the market yet.

What type of JOIN should we use here?

Answer: Left JOIN

**Query:**

```
SELECT *  
FROM farmers_market.customer AS c  
LEFT JOIN farmers_market.customer_purchases AS cp  
ON c.customer_id = cp.customer_id;
```

- There can be customers without any purchases.
- The customer table has details of all the customers regardless of their purchase history.
- Since we did a LEFT JOIN, we're getting a list of all customers, and their associated purchases, if there are any.
- Customers with multiple purchases will show up in the output multiple times for each item purchased.
- Customers without purchases will have NULL values in all fields displayed from the customer\_purchases table.

**To get the list of customers that did not purchase anything.**

**Query 1:**

```
SELECT c.* -- select columns from customer table only  
FROM customer AS c  
LEFT JOIN customer_purchases AS cp  
ON c.customer_id = cp.customer_id  
WHERE cp.customer_id IS NULL;
```

- Here we only selected columns from the customer table, using c.\*,
- because all of the columns on the customer\_purchases side of the relationship will be NULL (since we're filtering to NULL customer\_id, and there are no purchases in the customer\_purchases table without a customer\_id, since it is a required field).

Q2: Get all the customers who have deleted their account from the market.

What type of JOIN should we use here?

Answer: Right JOIN

**Query:**

```
SELECT *
FROM farmers_market.customer AS c
RIGHT JOIN farmers_market.customer_purchases AS cp
ON c.customer_id = cp.customer_id
```

**To get the list of customers who deleted their account.**

**Query 2:**

```
SELECT cp.* -- select columns from customer_purchases table
FROM customer AS c
RIGHT JOIN customer_purchases AS cp
ON c.customer_id = cp.customer_id
WHERE c.customer_id IS NULL;
```

- Here we only selected columns from the customer\_purchases table, using cp.\*,

- because all of the columns on the customer side of the relationship will be NULL (since we're filtering to NULL customer\_id, and there are no customers in the customers table without a customer\_id).

Now let's combine these two queries i.e. Query 1 & Query 2, to obtain the final solution query.

### Soln using UNION -

```
SELECT c.customer_id,  
       "New Customer" AS customer_type  
FROM farmers_market.customers AS c  
LEFT JOIN farmers_market.customer_purchases AS cp  
ON c.customer_id = cp.customer_id  
WHERE cp.customer_id IS NULL  
UNION DISTINCT  
SELECT cp.customer_id,  
       "Deleted Customer" AS customer_type  
FROM farmers_market.customers AS c  
RIGHT JOIN farmers_market.customer_purchases AS cp  
ON c.customer_id = cp.customer_id  
WHERE c.customer_id IS NULL;
```

### Output:

customer_id	customer_type
59	Deleted Customer
58	Deleted Customer
57	Deleted Customer
56	New Customer
55	New Customer



