

Databases · concept board

SQL & Relational Databases

Joins, aggregation, window functions, indexes, transactions, normalisation

1. Query order and basics

5. Indexes and performance

2. Joins

6. Transactions and isolation

3. Aggregation and subqueries

7. Schema design and normalisation


4. Window functions (top interview topic)

Draw on any slide, then Present (arrow keys) or export the board as a PDF.

1. Query order and basics

Definition

A relational database stores data in tables with a fixed schema, linked by keys, queried with SQL, and protected by ACID transactions (MySQL, PostgreSQL, SQL Server, SQLite).



Logical order SQL runs in — why WHERE can't use a SELECT alias

Sample tables used below (run me first)

```
DROP TABLE IF EXISTS emp;
DROP TABLE IF EXISTS dept;
CREATE TABLE dept (id INTEGER PRIMARY KEY, name TEXT NOT NULL);
CREATE TABLE emp (
  id INTEGER PRIMARY KEY,
  name TEXT NOT NULL,
  salary INTEGER NOT NULL,
  dept_id INTEGER REFERENCES dept(id),
  manager_id INTEGER REFERENCES emp(id)
);
INSERT INTO dept VALUES (1, 'Engineering'), (2, 'Sales'), (3, 'HR');
INSERT INTO emp VALUES
  (1, 'Asha', 150000, 1, NULL),
  (2, 'Ben', 120000, 1, 1),
  (3, 'Chen', 120000, 1, 1),
-- ... 4 more lines in the notes
```

- Primary key: unique, not null row id. Foreign key: references another table's key.
- WHERE filters rows before grouping; HAVING filters groups after.
- NULL means unknown: use IS NULL, never = NULL; COUNT(col) skips NULLs, COUNT(*) doesn't.

2. Joins

Join	Returns
INNER JOIN	only matching rows in both tables
LEFT JOIN	all rows from the left table, <i>NULL</i> s where there's no match
RIGHT JOIN	all rows from the right table
FULL OUTER JOIN	all rows from both
CROSS JOIN	every combination
SELF JOIN	a table joined to itself (employee → manager)

LEFT JOIN and SELF JOIN

-- uses the sample emp / dept tables

-- every employee with their department (Farah has none)

```
SELECT e.name, d.name AS dept
FROM emp e
LEFT JOIN dept d ON d.id = e.dept_id;
```

-- employees with their manager's name

```
SELECT e.name AS employee, m.name AS manager
FROM emp e
LEFT JOIN emp m ON m.id = e.manager_id;
```

-- departments with no employees (anti-join)

```
SELECT d.name FROM dept d
LEFT JOIN emp e ON e.dept_id = d.id
WHERE e.id IS NULL;
```

3. Aggregation and subqueries

GROUP BY, HAVING, subquery

-- uses the sample emp / dept tables

```
SELECT d.name, COUNT(e.id) AS people, AVG(e.salary) AS avg_salary
FROM dept d
JOIN emp e ON e.dept_id = d.id
GROUP BY d.id
HAVING COUNT(e.id) ≥ 2
ORDER BY avg_salary DESC;
```

-- earning more than the company average

```
SELECT name, salary FROM emp
WHERE salary > (SELECT AVG(salary) FROM emp);
```

4. Window functions (top interview topic)

Definition

Window functions compute across related rows without collapsing them (unlike GROUP BY): `fn() OVER (PARTITION BY ... ORDER BY ...)`.

Second-highest salary and top earner per department

```
-- uses the sample emp / dept tables

-- ROW_NUMBER vs RANK vs DENSE_RANK (Ben and Chen tie)
SELECT name, salary,
       ROW_NUMBER() OVER (ORDER BY salary DESC) AS row_num,
       RANK()        OVER (ORDER BY salary DESC) AS rnk,
       DENSE_RANK() OVER (ORDER BY salary DESC) AS dense
FROM emp;

-- 2nd highest distinct salary
SELECT DISTINCT salary FROM emp ORDER BY salary DESC LIMIT 1 OFFSET 1;

-- highest paid per department
WITH ranked AS (
  SELECT name, dept_id, salary,
  -- ... 4 more lines in the notes
```

- `LAG()` / `LEAD()` look at the previous / next row (month-over-month change).
- Running total: `SUM(amount) OVER (ORDER BY day)`.
- CTEs (`WITH ...`) name subqueries; recursive CTEs walk hierarchies.

5. Indexes and performance

Definition

An index (usually a B-tree) lets the database find rows without scanning the whole table — faster reads, slightly slower writes and extra storage.

Composite index and EXPLAIN

```
-- uses the sample emp / dept tables
CREATE INDEX idx_emp_dept_salary ON emp (dept_id, salary);

EXPLAIN QUERY PLAN
SELECT name FROM emp WHERE dept_id = 1 AND salary > 100000;
```

- Index columns used in WHERE, JOIN and ORDER BY; high-selectivity columns first.
- Composite index (a, b) helps WHERE a = ? and WHERE a = ? AND b > ?, but not WHERE b = ? alone (leftmost-prefix rule).
- Functions on a column (WHERE LOWER(email) = ...) skip a normal index.
- Use EXPLAIN / EXPLAIN ANALYZE to see scans vs index use; avoid SELECT *; fix N+1 queries with joins or batching.

6. Transactions and isolation

- ACID: Atomic (all or nothing), Consistent (constraints hold), Isolated (concurrent transactions don't interfere), Durable (committed survives crashes).

Isolation level	Prevents
Read uncommitted	nothing (dirty reads possible)
Read committed (Postgres default)	dirty reads
Repeatable read (MySQL InnoDB default)	+ non-repeatable reads
Serializable	+ phantom reads (safest, slowest)

Atomic transfer

```
DROP TABLE IF EXISTS account;  
CREATE TABLE account (id INTEGER PRIMARY KEY, balance INTEGER NOT NULL CHECK (balance ≥ 0));  
INSERT INTO account VALUES (1, 1000), (2, 200);
```

```
BEGIN;  
UPDATE account SET balance = balance - 500 WHERE id = 1;  
UPDATE account SET balance = balance + 500 WHERE id = 2;  
COMMIT;
```

```
SELECT * FROM account;
```

7. Schema design and normalisation

Normal form	Rule
1NF	atomic values, no repeating groups (no "tags: a,b,c" column)
2NF	1NF + every column depends on the whole primary key
3NF	2NF + no column depends on another non-key column

- Many-to-many needs a junction table (student_course(student_id, course_id)).
- Denormalise deliberately for read-heavy reporting (with care to keep copies in sync).
- Constraints keep data correct: NOT NULL, UNIQUE, CHECK, FOREIGN KEY ... ON DELETE CASCADE/RESTRICT.
- Soft delete (deleted_at) when you need history or undo.

Revise

FROM → WHERE → GROUP BY → HAVING → SELECT → ORDER BY. LEFT JOIN + IS NULL for "missing" rows. Window functions for rankings. Index WHERE/JOIN columns (leftmost prefix). ACID + isolation levels. Normalise to 3NF, then denormalise on purpose.

Rapid-fire Q&A (1/2)

WHERE vs HAVING?

WHERE filters rows before grouping;
HAVING filters groups after aggregation.

INNER vs LEFT JOIN?

INNER returns only matching rows;
LEFT returns every left-table row,
with NULLs where there is no match.

How do you find the second-highest salary?

```
SELECT DISTINCT salary ... ORDER BY  
salary DESC LIMIT 1 OFFSET 1, or  
DENSE_RANK() = 2.
```

ROW_NUMBER vs RANK vs DENSE_RANK

ROW_NUMBER is always unique; RANK
gives ties the same rank and skips
the next numbers; DENSE_RANK gives
ties the same rank without gaps.

What is an index and its trade-off?

A data structure (usually a B-tree)
that speeds up lookups; it costs
storage and makes writes slower.

What is the leftmost-prefix rule?

A composite index (a, b, c) is used
for filters on a, a+b, or a+b+c — not
for b or c alone.

Rapid-fire Q&A (2/2)

Explain ACID.

Atomicity, Consistency, Isolation, Durability — the guarantees of a database transaction.

What is normalisation?

Organising tables to remove duplication and update anomalies: 1NF atomic values, 2NF full key dependency, 3NF no transitive dependency.

DELETE vs TRUNCATE vs DROP?

DELETE removes chosen rows (logged, can roll back); TRUNCATE quickly removes all rows; DROP removes the table itself.

What is the N+1 query problem?

One query for a list, then one more per item. Fix with a JOIN, an IN (...) batch, or eager loading.

UNION vs UNION ALL?

UNION removes duplicates (extra work); UNION ALL keeps all rows.