



SQL Performance Checklist

A checklist to help you investigate and improve the performance of an SQL query..

How to Use:

Go through each of the checklist items and answer Yes, No, or Don't Know.

Use the answers to determine any next steps or areas to improve.

Section 1: Query Context and Expectations

- Is there a clear target runtime for this query that stakeholders agree on?
- Is the query slow relative to its historical performance, not just “feels slow”?
- Does the query run frequently enough that small delays add up?
- Is the slowness visible to users or downstream systems?
- Is the query executed as part of an important workflow?

Section 2: Data Characteristics

- Has the underlying data volume grown significantly since the query was written?
- Are a small number of values responsible for a large portion of the rows?
- Does the query operate on tables that grow continuously?
- Is the query sensitive to date ranges or rolling windows?
- Does performance vary depending on which values are queried?



Section 3: Query Shape & Logic

- Does the query join multiple large tables together?
- Does the query include aggregations over large result sets?
- Are there nested queries or derived result sets involved?
- Does the query return more columns than are actually needed?
- Is the query reused in multiple places with different inputs?

Section 4: Execution Plan Signals

- Is one step responsible for the majority of the total execution cost?
- Are estimated row counts significantly different from actual rows processed?
- Does the plan show work being repeated for many rows?
- Are large data sets being processed earlier than expected?
- Does the execution plan change depending on input values?
- Does the plan include warnings or flagged operators?



Section 5: Index & Access Path Signals

- Is a large amount of data being read compared to rows returned?
- Does the plan show full data access where targeted access was expected?
- Is filtering happening after data is read rather than during access?
- Are multiple access paths combined to satisfy the query?
- Does performance get worse as data grows?
- Does the same query use different access paths at different times?

Section 6: Environment & Usage

- Does the query behave differently in production compared to test or dev?
- Is the query executed concurrently by many users or processes?
- Is the query run interactively and also by background jobs?
- Does performance vary depending on time of day or system load?
- Is the query sensitive to parameter values or user input?