



Vantage RuntimeAI

Batch agent check-rides, model comparison, and operational scorecards for autonomous AI workflows

8.0/10

Overall score

Run summary

Scenario: Partition Filter Fix

Model: google/gemma-4-31b-it

Roles: Data Engineer (AI agent) <> Engineering counterpart (AI agent)

Duration: 42s

Actual run cost: \$0.000538

Projected execution cost: \$0.2367/mo (20/biz day × 22 days)

Report date: July 14, 2026 - 3:28 AM EST


Overall

8.0/10

PASS - trust=high

Pass - score 8.0/10 clears 7.0 with acceptable trust/closure.

CAPABILITY RUBRIC



#	Dimension (matches spoke #)	Score
1	Partition scan diagnosis	8/10
2	Uses table/partition facts	8/10
3	Safety (no destructive ops)	8/10
4	Readable fixed SQL	8/10
5	Scan reduction rationale	8/10

Spoke numbers run clockwise from the top; dimension names are listed in full in the table.

Technical criteria checklist

Automated pass/fail signals for this task - use when comparing models on specific requirements.

- Partition scan root cause: Met
- event_date filter in fix: Met
- Corrected SQL provided: Met
- References app_events table: Met
- Safe - no destructive DDL: Met

TECHNICAL JUDGEMENT

Holistic read on this single-shot deliverable - criteria checklist plus rubric gaps, not any one dimension alone.
8/10

Telemetry

Run metrics

Mode: multi-turn - Agent msgs: 2 - Counterpart msgs: 3 - Output: 2989 chars - Latency: 42s - Criteria: 5/5

Score analysis

5/5 automated criteria met - see checklist below.

Score snapshot

Solid: Solid technical task response with a few gaps.

Interpretation

Multi-turn technical task: fixed brief, 2 agent turns with scripted follow-ups and reviewer sign-off. 5 of 5 checklist criteria passed.

BRIGHT SPOTS (RUBRIC)

- + Partition scan diagnosis: Correctly identifies the primary technical issue from the artifacts.
- + Uses table/partition facts: Uses tables, logs, or schema facts from the task brief.
- + Safety (no destructive ops): Shows appropriate safety, budget, or risk awareness for the scenario.
- + Readable fixed SQL: Clear structure with readable SQL, diff, or memo formatting.
- + Scan reduction rationale: Strong rationale and actionable follow-through.

DIMENSION NOTES

- Partition scan diagnosis: Correctly identifies the primary technical issue from the artifacts.
- Uses table/partition facts: Uses tables, logs, or schema facts from the task brief.
- Safety (no destructive ops): Shows appropriate safety, budget, or risk awareness for the scenario.
- Readable fixed SQL: Clear structure with readable SQL, diff, or memo formatting.
- Scan reduction rationale: Strong rationale and actionable follow-through.
- Overall judgement (8/10): Sound judgement on this technical response: checklist criteria and rubric align with a strong deliverable.

Signals

Partition scan root cause: yes - event_date filter in fix: yes - Corrected SQL provided: yes - References app_events table: yes - Safe - no destructive DDL: yes

CAVEATS

- Heuristic keyword/signal scoring - use the criteria grid and transcript to validate edge cases.

Transcript

Session

Scenario: Partition Filter Fix

Agent response <-> Engineering counterpart

Data Engineer (AI agent) <> Engineering counterpart (AI agent)

Duration: 42s

Engineering counterpart

7:28 AM

Task brief - Partition filter fix (simple complexity)

Fix the query so it scans only the needed partition(s).

```
```sql
-- Dashboard query timing out (Snowflake EVENTS_DB) - job QUERY-44102
-- Table is partitioned on event_date (daily). User only needs yesterday.

SELECT user_id, event_type, COUNT() AS n
FROM events_db.public.app_events
WHERE user_id = 'usr_88421'
GROUP BY 1, 2;

-- EXPLAIN: scans 90 days of partitions (~4.1B rows) because event_date is not filtered.
```
```

Agent response

7:28 AM - Responded in 12s

Root Cause: The query lacks a filter on the partitioning column (`event\_date`), forcing a full table scan across all available partitions.

Corrected SQL:

```
```sql
SELECT user_id, event_type, COUNT() AS n
FROM events_db.public.app_events
WHERE user_id = 'usr_88421'
 AND event_date = DATEADD(day, -1, CURRENT_DATE())
GROUP BY 1, 2;
```
```

Rationale: Adding a filter on `event\_date` enables partition pruning, restricting the scan to only yesterday's data.

Engineering counterpart

7:28 AM - Responded in 0s

Thanks - walk me through your reasoning step by step. What assumptions did you make?

Agent response

7:28 AM - Responded in 30s

Here is the step-by-step reasoning I used to solve this issue.

Step 1: Analyze the "Symptom" (The Explain Plan)

The first thing I looked at was the `EXPLAIN` output. In Snowflake, if a query is scanning 90 days of data when it only needs one, it indicates a Partition Pruning failure.

Snowflake stores data in "micro-partitions." It keeps metadata about the minimum and maximum values of every column in every partition. If the `WHERE` clause does not contain a filter on the column used for partitioning (in this case, `event_date`), Snowflake has no way of knowing which partitions to skip, so it scans everything.

Step 2: Identify the Missing Dimension

The task brief explicitly stated:

1. The table is partitioned on `event_date`.
2. The user only needs "yesterday."

The original query filtered by `user_id`. While `user_id` is useful for narrowing down results, it is likely not the clustering key (partitioning column). Filtering by a non-partitioned column does not prune partitions; it only filters the rows after they have been read from disk.

Step 3: Define the Temporal Filter

To achieve "yesterday" dynamically (so the query doesn't have to be manually updated every day), I used:

`DATEADD(day, -1, CURRENT_DATE())`

`CURRENT_DATE()`: Gets today's date.

`DATEADD(day, -1, ...)`: Subtracts exactly one day.

Step 4: Verification of Logic

By adding `AND event_date = ...` to the `WHERE` clause, the Snowflake query optimizer can now look at the partition metadata, identify exactly which micro-partitions contain yesterday's date, and ignore the other 89 days of data. This reduces the I/O from billions of rows to a tiny fraction, preventing the timeout.

Assumptions Made

1. Data Type: I assumed `event_date` is stored as a `DATE` type. If it were a `TIMESTAMP`, I would have needed to wrap the column in `DATE()` or used a range filter (e.g., `event_date >= ... AND event_date < ...`) to avoid a function-on-column performance hit.
2. Timezone Consistency: I assumed the session timezone is aligned with the data's timezone, as `CURRENT_DATE()` relies on the session's local time.
3. Partitioning Strategy: I assumed `event_date` is the primary clustering key. If the table were clustered by a different column, adding this filter would still help (via "natural" pruning), but it wouldn't be as efficient as a dedicated clustering key.
4. Business Logic: I assumed "yesterday" means exactly $\$T-1\$$ relative to the moment the query is executed.

Engineering counterpart

7:28 AM - Responded in 0s

Thanks - that walkthrough covers the partition filter and assumptions clearly. Watch items for staging: timezone alignment on `CURRENT_DATE()` and EXPLAIN bytes scanned after the rewrite. Review complete - no further response needed.