



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/gemini-2.5-flash

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

Duration: 5s

Actual run cost: \$0.002650

Projected execution cost: \$1.17/mo (20/biz day × 22 days)

Report date: July 13, 2026 - 11:23 PM 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: 3188 chars - Latency: 5s - 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: 5s

Engineering counterpart

3:23 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

3:23 AM - Responded in 1s

Root Cause: The query is missing a filter on the `event\_date` column, leading to a full partition scan.

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

Rationale: Filtering on `event\_date` directly utilizes the table's partitioning strategy.

Engineering counterpart

3:23 AM - Responded in 0s

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

Agent response

3:23 AM - Responded in 4s

Let's break down the thought process step-by-step.

1. Understanding the Problem:

Goal: Fix a Snowflake query that's timing out due to scanning too many partitions.

Context: The table `events\_db.public.app\_events` is partitioned by `event\_date` (daily).

User Need: The user "only needs yesterday's" data.

Current Query:

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

Symptom: `EXPLAIN` shows it scans 90 days of partitions (~4.1B rows). This confirms a "partition-scan issue."

2. Identifying the Core Issue (Root Cause):

The `EXPLAIN` output explicitly states "event_date is not filtered."

The table is partitioned by `event\_date`.

If we don't filter on the partition key, the database system (Snowflake in this case) has no way of knowing which partitions not to scan. It has to look at every partition to find all `user\_id = 'usr\_88421'`, even if that user's data for a specific day might only reside in one partition.

Conclusion: The query is missing a filter on `event\_date`.

3. Determining the Required Filter Condition:

The problem statement explicitly says: "User only needs yesterday."

How do we get "yesterday" in SQL?

`CURRENT\_DATE` gives today's date.

Subtracting one day from `CURRENT\_DATE` gives yesterday's date.

Standard SQL `INTERVAL` syntax is `INTERVAL '1 day'` (or `INTERVAL '1' day`).

So, `CURRENT\_DATE - INTERVAL '1 day'` is the correct expression for yesterday.

4. Constructing the Corrected SQL:

Take the original query.

Add the `event\_date` filter to the `WHERE` clause using an `AND` operator (since `user\_id` filter already exists).

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

5. Formulating the Rationale:

The goal is to reduce partition scans.

The table is partitioned on `event\_date`.

Therefore, filtering on `event\_date` will allow Snowflake to perform "partition pruning" (or "scan pruning"). It will only scan the specific partition(s) that match the `event\_date` condition.

Concise Rationale: Filtering on `event\_date` directly utilizes the table's partitioning strategy.

6. Reviewing Constraints and Deliverables:

- (1) Root cause in one sentence: "The query is missing a filter on the `event\_date` column, leading to a full partition scan."
- Check.
 - (2) Corrected SQL with appropriate event_date filter: Done. - Check.
 - (3) One-line rationale: Done. - Check.
- Do not suggest

Engineering counterpart**3:23 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.