



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-3-12b-it

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

Duration: 15s

Actual run cost: \$0.000255

Projected execution cost: \$0.1122/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: 3727 chars - Latency: 15s - 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: 15s

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 2s

(1) The query lacks a filter on the `event\_date` column, causing it to scan an excessively large number of partitions.

(2)

```
```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;
```
```

(3) Filtering by yesterday's date ensures the query only scans the relevant partition, significantly reducing scan size and improving performance.

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 13s

Okay, let's break down the reasoning step-by-step. Here's the thought process behind the solution, including the assumptions I made:

1. Understanding the Problem

The Situation: A dashboard query is timing out in Snowflake. The `EXPLAIN` plan reveals the query is scanning 90 days of partitions, processing roughly 4.1 billion rows. This is clearly inefficient and the root cause of the timeout.

Key Information:

The table `events\_db.public.app\_events` is partitioned by `event\_date` (daily). This means Snowflake divides the data into separate partitions based on the date of the event.

The user only needs data from "yesterday." This implies a specific, relatively narrow time window of interest.

The query already filters on `user\_id = 'usr\_88421'`, which is good for narrowing results but isn't enough to address the partition scan issue.

2. Identifying the Root Cause (and my assumption)

The problem is not the `user\_id` filter. That's helpful, but the real performance killer is the lack of a date filter. Without it, Snowflake has to scan potentially all partitions.

Assumption: The prompt states the user "only needs yesterday's data." I assumed they wanted the data from the day prior to the current execution date. This is a very standard interpretation of "yesterday" in a database context. This is a crucial assumption-if the requirement was something else (like a specific date range, or data from a fixed point in the past), the solution would change.

3. Crafting the Solution

Goal: Add a filter to the `WHERE` clause that targets only the partition containing yesterday's data.

SQL Approach:

`event_date = CURRENT_DATE() - INTERVAL '1' DAY`

`CURRENT_DATE()`: Returns the current date.

`INTERVAL '1' DAY`: Represents an interval of one day. Subtracting this from `CURRENT_DATE()` gives us yesterday's date.

Combining them, the filter `event_date = CURRENT_DATE() - INTERVAL '1' DAY` isolates the partition for yesterday's events.

Complete Query: Integrate the filter into the existing query:

```
```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;
```
```

4. Rationale

Why it works: By filtering on `event\_date`, we instruct Snowflake to only read the partition corresponding to yesterday. This dramatically reduces the amount of data scanned, leading to faster query execution and preventing timeouts.

One-line rationale: Filtering by yesterday's date ensures the query only scans the relevant partition, significantly reducing scan size and improving performance.

Important Considerations & Alternatives (That I didn't pursue based on the prompt)

Parameterization: In a production environment, it's generally better to parameterize the date rather than hardcoding it.

This allows the query to be reused with different dates. For example, you could pass in the date as an argument to a stored procedure.

Time Zones: `CURRENT_DATE()` is time zone dependent. If the `event_date` column does not store dates in the same time zone as the

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.