



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: deepseek/deepseek-chat

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

Duration: 16s

Actual run cost: \$0.001492

Projected execution cost: \$0.6566/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: 2563 chars - Latency: 16s - 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: 16s

Engineering counterpart

3:22 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

3:22 AM - Responded in 4s

Root cause: The query scans 90 days of partitions because it lacks a filter on the `event\_date` column, which is the table's partitioning key.

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 = CURRENT_DATE() - 1
GROUP BY 1, 2;
```
```

Rationale: Adding `event\_date = CURRENT\_DATE() - 1` ensures the query scans only yesterday's partition, significantly reducing the scan volume.

Engineering counterpart

3:22 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 12s

Certainly! Here's a step-by-step breakdown of my reasoning and the assumptions I made:

Step 1: Understand the Problem

- The query is timing out because it scans 90 days of partitions, resulting in ~4.1B rows being processed.
- The table is partitioned on `event\_date` (daily), and the user only needs data from yesterday.

Step 2: Analyze the Query

- The query filters on `user\_id` but does not filter on `event\_date`.
- Without a filter on the partitioning key (`event\_date`), Snowflake scans all partitions by default, which is inefficient.

Step 3: Identify the Fix

- To limit the scan to only yesterday's data, we need to add a filter on `event\_date`.
- Since the user only needs yesterday's data, we use `event\_date = CURRENT\_DATE() - 1`.

Step 4: Write the Corrected Query

- Add the `event\_date` filter to ensure the query scans only the relevant partition.

```
```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() - 1
GROUP BY 1, 2;
```
```

Step 5: Validate Assumptions

1. Assumption 1: The table is partitioned on `event\_date`, and the user only needs data from yesterday.
 - Confirmed by the task description.
2. Assumption 2: The `event\_date` column contains daily dates in a format compatible with `CURRENT\_DATE()`.
 - If the column uses a different format (e.g., timestamp or string), the filter would need adjustment.
3. Assumption 3: Adding a filter on `event\_date` will reduce the scan volume to only yesterday's partition.
 - Correct, as Snowflake's partition pruning will exclude irrelevant partitions.

Step 6: Rationale

Adding `event\_date = CURRENT\_DATE() - 1` ensures the query scans only yesterday's partition, reducing the scan volume from 90 days to 1 day and improving performance. This is a non-destructive fix that doesn't require any DDL changes.

Let me know if you'd like to adjust the assumptions or explore further!

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.