

Snowflake Optimization Pack

Discovery questionnaire — score each item honestly; low scores are the point: they show where the program will hurt.

Pack	Snowflake Optimization Pack
Document	Discovery questionnaire (24 questions)
Response scales	Maturity 1–5 · Yes / Partially / No · Written
Date	2026-09

AnovaCloud — Engineering the AI-Native Enterprise. Frisco, Texas. This pack is a working instrument, not a brochure: complete it with your team, score honestly, and bring it to your discovery call. © 2026 AnovaCloud.

HOW TO USE

Run this like a workshop, not a survey

Complete it with the stakeholders named in each section's intro. One scribe, one room (or call), 90 minutes. Score from evidence — monitoring data, documents, demos — not opinions. Any item scoring 1–2 is a program risk: it needs an owner and a close date before you commit budget. Bring the scored questionnaire to your discovery call; it compresses weeks of fact-finding into one session.

1 · WAREHOUSES & WORKLOADS**Warehouses & workloads**

Warehouses are the unit of cost and performance.

Q1. How many warehouses exist, and is each mapped to a workload with an owner?

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: warehouse inventory with owners.

Q2. Are sizes right (not XL-by-default), with auto-suspend tuned per workload?

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: utilization vs. size analysis.

Q3. Is concurrency handled by multi-cluster/scale policy — or by queuing?

Evidence: queuing metrics during peak.

Q4. Are workloads isolated (EL vs. BI vs. ad-hoc) or sharing warehouses and blaming each other?

Yes

Partially

No

Evidence: workload-to-warehouse mapping.

2 · MODELING & DATA DESIGN**Modeling & data design**

The cheapest query is the one that scans less.

Q5. Are clustering keys defined on large tables with selective filters?

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: clustering depth + query pruning stats.

Q6. Are fact/dimension models sane, or are analysts joining raw event tables directly?

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: model review of top schemas.

Q7. Is result caching actually hit (or is every dashboard a full recompute)?

Evidence: cache-hit ratio on top dashboards.

Q8. Are materialized/dynamic tables used where recompute is wasteful?

Yes

Partially

No

Evidence: inventory of pre-aggregation.

3 · COST & CREDIT BURN

Cost & credit burn

You can't optimize what you can't attribute.

Q9. What are the top-10 credit consumers (warehouse, query, user, team)?

Evidence: ACCOUNT_USAGE top-consumer query.

Q10. Is credit burn attributed to teams with budgets and alerts?

1	2	3	4	5
---	---	---	---	---

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: attribution dashboards.

Q11. What share of spend is idle (warehouses running with no queries)?

Evidence: idle-credit measurement.

Q12. Are storage costs (incl. Time Travel/Failsafe retention) reviewed, not just compute?

☐	Yes	☐	Partially	☐	No
--------------------------	-----	--------------------------	-----------	--------------------------	----

Evidence: storage breakdown.

4 · PERFORMANCE & CONCURRENCY

Performance & concurrency

Slow dashboards kill trust; queuing kills mornings.

Q13. What is p95 dashboard/query latency for the top consumer workloads?

Evidence: measured p95s.

Q14. Are the slowest queries profiled (Query Profile: pruning, spilling, join explosion)?

1	2	3	4	5
---	---	---	---	---

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: profile notes on top-20 queries.

Q15. Is search optimization / pruning effective on selective large-table queries?

☐	Yes	☐	Partially	☐	No
--------------------------	-----	--------------------------	-----------	--------------------------	----

Evidence: before/after on targeted queries.

Q16. Do peak periods queue, and is the queuing costed (user wait + credits)?

Evidence: queuing incidents + impact.

5 · GOVERNANCE & SECURITY**Governance & security**

Credits aren't the only thing that can leak.

Q17. Are RBAC roles least-privilege with managed access schemas?

☐ 1	☐ 2	☐ 3	☐ 4	☐ 5
----------------------------	----------------------------	----------------------------	----------------------------	----------------------------

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: role review sample.

Q18. Is PII classified with masking policies and row access where required?

☐ Yes	☐ Partially	☐ No
------------------------------	------------------------------------	-----------------------------

Evidence: classification + policy inventory.

Q19. Are data sharing / marketplace exposures inventoried and approved?

☐ 1	☐ 2	☐ 3	☐ 4	☐ 5
----------------------------	----------------------------	----------------------------	----------------------------	----------------------------

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: share inventory.

Q20. Is query history monitored for anomalous access (not just cost)?

☐ Yes	☐ Partially	☐ No
------------------------------	------------------------------------	-----------------------------

Evidence: monitoring rules.

6 · OPERATIONS**Operations**

Snowflake still needs operating.

Q21. Are resource monitors with alert + suspend actions configured per workload?

☐ Yes ☐ Partially ☐ No

Evidence: resource monitor list.

Q22. Is there a query-tagging standard so every credit is attributable?

☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: tagging compliance %.

Q23. Are dbt/other transformations version-controlled with CI and tests?

☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: repo + CI + test coverage.

Q24. Is there a runbook for cost spikes and performance incidents?

☐ Yes ☐ Partially ☐ No

Evidence: runbook link + drill record.

SCORING

Read the pattern, not just the average

Average each section. Sections averaging below 3 are your critical path — sequence the project plan so those risks close first. A single 1 on a load-bearing question (downtime budget, rollback, access control) outweighs a 4 average elsewhere. Record owners and close dates below.

Lowest-scoring section:

Single biggest risk:

Owner:

Close date:
