DATALYTIQ · FREE GUIDE

Data Interview Prep Kit

The core of what gets tested in analytics and data science interviews, and what the interviewer is actually listening for. No list of 200 questions.

By Eromonsele Oboh · datalytiq.ca · November 2026

Which roles this fits

Built for **Data Analyst**, **Product Analyst** and **Product Data Scientist** interviews, in tech, media and similar companies. It's not built for ML Engineer, Research Scientist or Data Engineer roles, which test modelling, coding and systems far more.

Section	Data Analyst	Product Analyst	Product Data Scientist
2. SQL	Core	Core	Core
3. Product and metrics	Sometimes	Core	Core
4. Experiments	Light	Often	Core
5. Stakeholder round	Senior roles	Core	Core
6. Your story	Core	Core	Core

1. What each round is really testing

- **Recruiter screen:** can you explain what you did and why it mattered, in plain words, in two minutes?
- **SQL:** clean, correct logic under time pressure. Talking through your approach counts as much as the final query.
- **Product and metrics:** can you pick the right number, and say what would make you doubt it?
- **Experiments:** do you know when a result can't be trusted?
- **Stakeholder round:** would a PM or engineering lead want you in the room when a decision gets made?
- **Behavioural:** have you done this before, and do you learn from what went wrong?

2. SQL ALL THREE ROLES

Three patterns come up again and again.

Cohort retention

"For each signup week, what share of users came back the following week?"

Listening for: you define "came back" before writing anything, you keep users with no activity (LEFT JOIN, not INNER JOIN), and you get the denominator right.

Did X but not Y

"Users who bought in January but not in February."

Listening for: `NOT EXISTS` or `LEFT JOIN ... IS NULL`, and knowing that `NOT IN` returns nothing if the list contains a NULL.

Running totals and ranking

"Monthly revenue with a running total", or "each user's first purchase".

Listening for: window functions (`SUM() OVER`, `ROW_NUMBER()`) with the right `PARTITION BY` and `ORDER BY`.

Common traps: joins that create duplicate rows and inflate counts, dates and time zones, and integer division turning your rate into 0.

3. Product and metrics

PRODUCT ANALYST · PRODUCT DS

"Daily active users dropped 10% overnight. What do you check?"

Strong order: Is it real, or a tracking change, outage or broken pipeline? Is it everywhere, or one platform, country, app version or user type? Did something ship, or happen outside (a holiday, an app store issue)? Only then dig into causes.

"How would you know if a new feature is working?"

Strong shape: start with what the feature is for. Pick one main metric tied to that goal, one or two guardrails that shouldn't get worse, and say how long you'd wait before judging.

"Pick one north star metric for this product."

Listening for: a number that reflects real value to the user, not just activity, and you say what it misses.

The pattern for all three: goal, metric, guardrail, caveat.

4. Experiments

PRODUCT DS CORE · PRODUCT ANALYST OFTEN

"Revenue is up 2% but retention is down 1%. Do you ship?"

Listening for: there's no single right answer. They want to hear you check that both changes are real, ask which matters more over a year, look for a segment driving the drop, and consider a partial launch or a fix and retest. Then you make a call.

"The test looked significant on day 3. Why not stop it?"

Listening for: checking repeatedly and stopping at the first good result raises false positives. Set the sample size and duration up front, and run at least a full week to cover weekday patterns.

Also know, in one plain sentence each: what a p-value means, what a sample ratio mismatch is (the split isn't what you set, so something is broken), and why new features can get a short novelty bump.

5. The stakeholder round SENIOR ROLES

It's easy to under-prepare for this round because it doesn't feel technical. At senior level it decides a lot.

What they're listening for:

- You turn analysis into a recommendation, not a list of findings.
- You push back with evidence, without blocking the work.
- You explain tradeoffs and risk in plain language.
- You ask what the other person needs before you answer.

Practice these out loud:

- "The PM wants to launch Friday. Your data says the feature hurts retention for new users. What do you say in the meeting?"
- "An engineer says your metric is wrong because the logging can't be trusted. How do you handle it?"

A good answer shape: acknowledge their goal, show the evidence briefly, offer a way forward (a smaller launch, a fix, or a date for more data).

6. Your story ALL THREE ROLES

Have one project you can tell well in two minutes. It will carry your behavioural, product and impact questions.

1. **Problem:** what was going wrong, and why it mattered to the business.
2. **What you did:** the analysis in one or two sentences. Skip the tooling unless asked.
3. **What you found:** one clear finding.
4. **Decision:** what changed because of it.
5. **Result:** a number if you have one. If not, what happened next.

Then prepare for the follow ups: "What would you do differently?", "How do you know it caused the change?", "Who disagreed with you?"

7. Three mistakes to avoid

1. Starting the query or the answer before asking a single clarifying question.
2. Giving one metric with no guardrail or caveat.
3. Ending without the "so what": what you'd recommend, and why.

Want a mock interview or feedback on your story? I mentor students and early career data professionals one on one. Book a call at **datalytiq.ca/mentorship**.