Data Team Hiring Playbook

Interview guides, skills rubrics & org design for data teams

ABOUT THIS PLAYBOOK

This playbook provides structured interview guides, skills rubrics, and org design templates for hiring data engineers, data scientists, analytics engineers, and ML engineers. Each role section includes competency expectations by level, interview questions with evaluation criteria, and red/green flags.

ROLES COVERED IN THIS PLAYBOOK

- **Data Engineer:** Builds and maintains data pipelines, data lakes, and streaming infrastructure.
- **Analytics Engineer:** Transforms raw data into trusted datasets using dbt; bridges DE and analyst.
- **Data Scientist:** Builds predictive models; communicates insights to business stakeholders.
- **ML Engineer:** Productionises ML models; owns MLOps pipelines, serving, and monitoring.

UNIVERSAL HIRING PRINCIPLES

- Structure every interview against a consistent rubric -- eliminate recency bias.
- Use real-world case studies, not abstract puzzles, to assess problem-solving.
- Evaluate communication skills explicitly -- data roles require translating complexity.
- Assess data intuition: can the candidate spot anomalies, question assumptions, and sanity-check results?
- Prioritise learning velocity over current tool knowledge -- stacks change; curiosity doesn't.
- Conduct a values/culture-add interview as a separate, scored stage.
- Require a structured debrief within 24h to prevent anchoring on first impressions.

LEVELLING FRAMEWORK

- **L3 - Junior:** 0-2 yrs. Executes defined tasks; learns from seniors; limited scope.
- **L4 - Mid:** 2-5 yrs. Owns features independently; mentors juniors; proactive.
- **L5 - Senior:** 5-8 yrs. Leads technical design; influences team direction; cross-functional.
- **L6 - Staff:** 8+ yrs. Org-wide technical strategy; drives standards; executive credibility.

Data Team Hiring Playbook

Data Engineer | Analytics Engineer

DATA ENGINEER -- SKILLS RUBRIC

- **Pipeline Design:** Can architect idempotent, observable batch and streaming pipelines.
- **SQL & Python:** Advanced SQL; Python for ETL, orchestration, and testing.
- **Cloud Platforms:** AWS/GCP/Azure data services; Terraform or CDK for IaC.
- **Orchestration:** Airflow, Prefect, or Dagster DAG design; dependency management.
- **Data Modelling:** Understands Kimball, data vault, and lakehouse patterns.
- **Observability:** Implements data quality checks, SLAs, and alerting.

DATA ENGINEER -- INTERVIEW QUESTIONS

- Walk me through a data pipeline you built end-to-end. What would you do differently now?
- How do you handle schema evolution in a production pipeline with downstream consumers?
- Design a real-time pipeline that processes 1M events/hour with exactly-once semantics.
- How do you test data pipelines? What makes a good data quality check?
- Describe a production data incident you caused or resolved. What did you learn?

ANALYTICS ENGINEER -- SKILLS RUBRIC

- **dbt:** Models, tests, snapshots, macros, and sources. Packages and DRY principles.
- **Data Modelling:** Fact/dimension design; slowly changing dimensions; grain definition.
- **SQL Mastery:** Window functions, CTEs, query optimisation, and EXPLAIN plan reading.
- **Business Acumen:** Understands P&L, funnel metrics, cohort analysis, and attribution.
- **Documentation:** YAML schema files, dbt docs, data dictionary ownership.
- **Stakeholder Mgmt:** Can say no, negotiate definitions, and manage metric disputes.

ANALYTICS ENGINEER -- INTERVIEW QUESTIONS

- How do you decide between a dbt model and a dashboard calculation? What are the tradeoffs?
- Walk me through how you'd model a subscription SaaS business in dbt from raw Stripe data.
- A business stakeholder says two dashboards show different revenue numbers. How do you resolve it?
- What's your approach to dbt testing? Which tests do you always add?

Data Team Hiring Playbook

Data Scientist | ML Engineer

DATA SCIENTIST -- SKILLS RUBRIC

- **Statistics:** Hypothesis testing, distributions, A/B test design, Bayesian reasoning.
- **ML Algorithms:** Regression, trees, ensembles, neural nets; knows when to use each.
- **Python / R:** pandas, scikit-learn, PyTorch/TensorFlow, statsmodels.
- **Experimentation:** Causal inference, experiment design, power calculations, MDE.
- **Communication:** Can explain model results to non-technical executives clearly.
- **Business Framing:** Translates vague business problems into precise ML problem statements.

DATA SCIENTIST -- INTERVIEW QUESTIONS

- A model has 99% accuracy on our test set but fails in production. What's your diagnostic process?
- How would you design an A/B test for a new ML-powered recommendation feature? What are the risks?
- Explain precision/recall tradeoff to a VP of Sales who has no ML background.
- You have 3 months to build a churn model. Walk me through your approach from data to deployment.
- What are the limitations of using historical data to train models for forward-looking predictions?

ML ENGINEER -- SKILLS RUBRIC

- **MLOps Platform:** Kubeflow, MLflow, Vertex AI, or Databricks MLflow; experiment tracking.
- **Model Serving:** FastAPI, TorchServe, Triton, or cloud endpoints; latency/throughput SLAs.
- **CI/CD for ML:** GitHub Actions or Jenkins pipelines for model code; automated retraining.
- **Monitoring:** Data drift, concept drift, model performance degradation detection.
- **Software Engineering:** Clean code, testing (unit + integration), code review, documentation.
- **Infrastructure:** Docker, Kubernetes, GPU provisioning, cost management.

ML ENGINEER -- INTERVIEW QUESTIONS

- Describe the MLOps stack at your last company. What would you improve?
- How do you monitor a production model for degradation? What triggers a retrain?
- A model serving endpoint has p99 latency of 800ms vs a 200ms SLA. How do you diagnose and fix it?
- Walk me through how you'd implement a feature store for a team of 10 data scientists.

Data Team Hiring Playbook

Org Design Templates | Red & Green Flags

ORG DESIGN: TEAM STRUCTURES BY COMPANY STAGE

Seed / Series A (1-3 data people)

Hire a data generalist or analytics engineer first. Avoid premature specialisation. One person should cover SQL, dashboards, and basic pipelines.

Series B (4-10 data people)

Split data engineering from analytics. Add a data scientist if ML is strategic. Consider a Head of Data or Data Lead to set standards.

Series C+ (10+ data people)

Add ML Engineers, Analytics Engineers, and Data Platform Engineers. Embed analysts in product squads. Appoint a CDO or VP of Data.

EVALUATION SCORECARD (USE PER INTERVIEW)

Technical Skills	Score (1-5): ___ | Notes: ________________________
Problem Solving	Score (1-5): ___ | Notes: ________________________
Communication	Score (1-5): ___ | Notes: ________________________
Business Acumen	Score (1-5): ___ | Notes: ________________________
Culture Fit / Values	Score (1-5): ___ | Notes: ________________________
Learning Agility	Score (1-5): ___ | Notes: ________________________

RED FLAGS & GREEN FLAGS

RED FLAGS -- Proceed with caution

- Can't explain past project impact in business terms
- Blames data quality for all problems
- Only comfortable in one tool or language
- No questions for the interviewer
- Can't describe a failure or mistake
- Dismisses stakeholders as "non-technical"

GREEN FLAGS -- Strong signals

- Quantifies impact (revenue, cost, time saved)
- Has fixed upstream data quality issues proactively
- Has learned new stacks on the job multiple times
- Asks deep, specific questions about your data challenges
- Shares clear failure story with learnings
- Has changed a stakeholder's mind with data evidence